



(19)

Europäisches Patentamt

European Patent Office

Office européen des brevets



(11)

EP 0 923 049 A2

(12)

EUROPEAN PATENT APPLICATION

(43) Date of publication:
16.06.1999 Bulletin 1999/24(51) Int. Cl.⁶: G06T 3/40

(21) Application number: 98117551.6

(22) Date of filing: 16.09.1998

(84) Designated Contracting States:
AT BE CH CY DE DK ES FI FR GB GR IE IT LI LU
MC NL PT SEDesignated Extension States:
AL LT LV MK RO SI

(30) Priority: 12.12.1997 US 989961

(71) Applicant:
Hewlett-Packard Company
Palo Alto, CA 94303 (US)(72) Inventor: Vondran, Gary L., Jr.
Winchester, MA 01890 (US)(74) Representative:
Schoppe, Fritz, Dipl.-Ing. et al
Schoppe & Zimmermann
Patentanwälte
Postfach 71 08 67
81458 München (DE)

(54) Apparatus for tetrahedral and pruned tetrahedral interpolation

(57) New interpolation techniques allow improved efficiency and speed in performing color space conversions. A radial interpolation technique (FIG. 7-8) accomplishes an interpolation by generating successive sub-cubes (FIG. 6a-6e). A value of a vertex of the final sub-cube generated is used as the result of the interpolation. Sub-cubes (FIG. 6a-6e) are generated by averaging (FIG. 5) a selected vertex value with the vertex values of each of the remaining vertices. A pruned radial interpolation technique (FIG. 9-12, FIG. 38-39) employs a subset of the vertex values of the initially selected cube to generate the result of the interpolation, thereby improving upon the efficiency of the radial interpolation. A tetrahedral interpolation technique (FIG. 16-17) accomplishes an interpolation by generating successive sub-cubes (FIG. 14). A value of a vertex of the final sub-cube generated is used as the result of the interpolation. Sub-cubes are generated by applying a mathematical relationship which allows computation of sub-cube vertex values through a series of logical AND, logical OR and averaging operations. A pruned tetrahedral interpolation technique (FIG. 15-16, FIG. 18-20, FIG. 40-41) employs a subset of the vertex values of the initially selected cube to generate the result of the interpolation, thereby improving upon the efficiency of the tetrahedral interpolation.

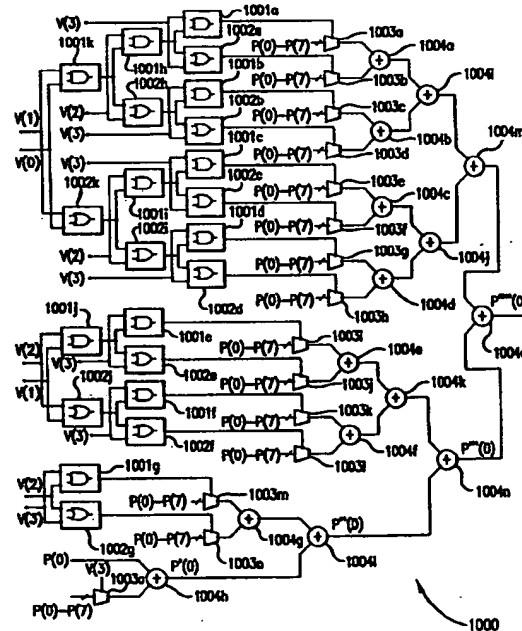


FIG. 18

Description**FIELD OF THE INVENTION**

- 5 [0001] This invention relates to the transformation of data, and more particularly to the transformation of data from a first space to a second space, such as in the conversion from a first color space to a second color space.

BACKGROUND OF THE INVENTION

- 10 [0002] Colorimetry has long been recognized as a complex science. In general, it has been found possible and convenient to represent color stimuli vectors in a three-dimensional space, called tristimulus space. Essentially, as defined in 1931 by the Commission Internationale L'Eclairage (CIE), three primary colors (X, Y, Z) can be combined to define all light sensations we experience with our eyes (that is, the color matching properties of an ideal trichromatic observer defined by specifying three independent functions of wavelength that are identified with the ideal observer's color matching functions form an international standard for specifying color). The fundamentals of such three-dimensional constructs are discussed in the literature, such as *Principles of Color Technology*, by Billmeyer and Saltzman, published by John Wiley & Sons, Inc., NY, copyright 1981 (2nd. ed.) and *Color Science: Concepts and Methods, Quantitative Data and Formulae*, by Wyszecki and Stiles, published John Wiley & Sons, Inc., copyright 1982 (2d ed.), incorporated herein by reference in pertinent parts, particularly pages 119 - 130.

- 20 [0003] Trichromatic model systems -- such as red, green, blue (RGB); cyan, magenta, yellow (CMY); hue, saturation, value (HSV); hue, lightness, saturation (HLS); luminance, red-yellow scale, green-blue scale ($L^a^*b^*$); luminance, red-green scale, yellow-blue scale (Luv); YIQ used in commercial color television broadcasting; and the like - provide alternatives for the system designer. See such works as *Fundamentals of Interactive Computer Graphics*, by Foley and Van Dam, Addison-Wesley Publishing Company, incorporated herein by reference in pertinent parts, particularly pages 606 - 621, describing a variety of tri-variable color models.

- 25 [0004] Color transformation between model systems in digital data processing presents may problems to the original equipment manufacturer. The translation of data from one system to another system is difficult because the relationship between the systems are generally non-linear. Therefore, a crucial problem is the maintaining of color integrity between an original image from an input device (such as a color scanner, CRT display, digital camera, computer software/firmware generation, and the like) and a translated copy at a output device (such as a CRT display, color laser printer, color ink-jet printer, and the like).

- 30 [0005] For example, computer artists want the ability to create a color image on a computer video and have a printer provide the same color in hard copy. Or, an original color photograph may be digitized with a scanner, resultant data may be transformed for display on a video monitor or reproduced as a hard copy by a laser, ink-jet or thermal transfer printer. As discussed in the reference materials cited, colors - be constructed as renderings of the additive primary colors, red, green and blue (RGB), or of the subtractive primary colors, cyan, magenta, yellow and black (CMYK). A transformation may require going from a RGB color space, for example, a computer video monitor, to a CMYK color space, for example, a laser printer hard copy. A transformation from one color space to another requires complex, non-linear computations in multiple dimensions. Some transform operations could be accomplished through matrix multiplication.

- 35 [0006] However, a difficulty in this method of color space conversion results from imperfections in the dyes, phosphors, and toners used for the production of the colors. An additional complication is that different types of media produce different color responses from printing with the same mixes of colorants. As a result, a purely mathematical color space conversion method does not provide acceptable color reproduction.

- 40 [0007] It has been recognized that superior results in color space conversion are obtained using a look up table scheme based upon a set of empirically derived values. Typically the RGB color space used for video displays use eight bits to represent each of the primary colors, red, green, and blue. Therefore, twenty four bits are required to represent each picture element. With this resolution, the RGB color space would consist of 2^{24} or 16,777,216 colors. Performing a color space conversion from each of these points in the RGB color space to generate the four CMYK (to maintain black color purity in printing, a separate black is usually provided rather than printing with all three of cyan, magenta, and yellow colorants to generate what is commonly known as process black) color space components would require a look-up table with 4×2^{24} or 67,108,864 bytes of data. The empirical construction of a look-up table with this number of entries is too costly.

- 45 [0008] In making the transform from one color space to another, a number of interpolation schemes well known in the field of color space conversion may be employed. Methods of performing color space conversion using trilinear interpolation, prism interpolation, and tetrahedral interpolation are disclosed in the published article PERFORMING COLOR SPACE CONVERSIONS WITH THREE DIMENSIONAL LINEAR INTERPOLATION, JOURNAL OF ELECTRONIC IMAGING, July 1995 Vol. 4(3), the disclosure of which is incorporated herein by reference. U.S. Patent No. 3, 893,166

(the disclosure of which is incorporated herein by reference), issued to Pugsley, discloses a scheme for translation between color spaces which uses a look-up table to access values used in an interpolation.

5 [0009] Conversion of large amounts of data between color spaces, such as is required for color printing, is a time consuming operation using the prior art methods of interpolation. The use of the computationally intensive prior art methods of interpolation for the color space conversion process makes high rates of data throughput difficult to achieve. A need exists for a interpolation method and interpolation apparatus that will enable a reduction in the computations required for performing a conversion between color spaces.

SUMMARY OF THE INVENTION

- 10 [0010] Accordingly, a pruned tetrahedral interpolator for interpolating between interpolation data values uses input data values each having d components to generate output data values. The d components are represented by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits. The d sets of higher order bits are used for selecting 2^d of the interpolation data values. The pruned tetrahedral interpolator includes a first multiplexer having a first multiplexer output and a first control input. The first multiplexer is configured for receiving the 2^d of the interpolation data values. The first control input is configured for receiving a first value determined from the d sets of lower order bits. The pruned tetrahedral interpolator further includes a first adder having a first input, a second input, and an output. The first input is coupled to the first multiplexer output and the second input is configured to receive one of the 2^d of the interpolation data values.
- 15 [0011] A tetrahedral interpolator for interpolating between interpolation data values uses input data values each having d components to generate output data values. The d components are represented, correspondingly, by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits. The d sets of higher order bits are used for selecting 2^d of the interpolation data values. The tetrahedral interpolator includes a set of 2×2^d multiplexers with each of the multiplexers having a multiplexer output and configured to receive the 2^d of the interpolation data values.
- 20 [0012] Each of the multiplexers is used for selecting a one of the 2^d of the interpolation data values responsive to one of 2×2^d values determined from the d sets of lower order bits. The tetrahedral interpolator further includes a set of 2^d adders with each of the adders having a first input, a second input, and an output. Each of a first group of 2^d of the multiplexers have the corresponding of the multiplexer output coupled to one of the first input. Each of a second group of 2^d of the multiplexers have the corresponding of the multiplexer output coupled to one of the second input. The set of 2×2^d multiplexers and the set of 2^d adders form a stage.
- 25 [0013] A method of tetrahedral interpolation uses interpolation data values for selection using input data values each having d components. The d components are represented by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits with each of the d sets of lower order bits having n bits. The d sets of lower order bits are designated as $1b_1, 1b_2, \dots, 1b_d$ with the bit position of each bit of the d sets of lower order bits designated from the most significant of the lower order bits to the least significant of the lower order bits by a value of i ranging, correspondingly, from n-1 to 0. The method of tetrahedral interpolation includes the step of computing a first value according to $v[i] = 2^{d-1} \times 1b_1[i] + 2^{d-2} \times b_2[i] + \dots + 2^{d-d} \times 1b_d[i]$ for the value of i equal to (n-1). The method of tetrahedral interpolation further includes the step of computing a first set of AND values according to $v[n-1] \& k$, for the value of k ranging from 2^{d-1} to 0. The method of tetrahedral interpolation further includes the step of computing a first set of OR values according to $v[n-1] | k$, for the value of k ranging from $2^d - 1$ to 0. The method of tetrahedral interpolation further includes the step of selecting a first set of 2^d pairs of the interpolation data values using the first set of the AND values and the first set of the OR values. Each of the first set of 2^d pairs are selected using one of the first set of the AND values and one of the first set of the OR values each computed using the same of the value of k. The method of tetrahedral interpolation further includes the step of computing a first set of 2^d sums by summing each of the first set of 2^d pairs of the interpolation data values.
- 30 [0014] A method of pruned tetrahedral interpolation uses interpolation data values for selection using input data values each having d components. The d components are represented by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits with each of the d sets of lower order bits having n bits. The d sets of lower order bits are designated as $1b_1, 1b_2, \dots, 1b_d$ with the bit position of each bit of the d sets of lower order bits designated from the most significant of the lower order bits to the least significant of the lower order bits by a value of i ranging, correspondingly, from n-1 to 0. The method of pruned tetrahedral interpolation includes a step of computing a first set of $2^n - 2$ values using bitwise AND operations and bitwise OR operations operating upon $v[i]$. Where $v[i]$ is equal to $2^{d-1} \times 1b_1[i] + 2^{d-2} \times b_2[i] + \dots + 2^{d-d} \times 1b_d[i]$ for the value of i ranging from (n-1) to 0. The method of pruned tetrahedral interpolation further includes a step of selecting at least the minimum of 2^n and 2^d of the interpolation data values using ones of the first set of $2^n - 2$ values, $v[i]$ for one of the values of i equal n-1, and the d sets of higher order bits. The method of pruned tetrahedral interpolation further includes a step of adding a second set of the interpolation data values formed from the interpolation data values from the step of selecting to generate a sum.
- 35 [0015] A pruned tetrahedral interpolator for interpolating between interpolation data values uses input data values each having d components. The d components are represented by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits with each of the d sets of lower order bits having n bits. The d sets of lower order bits are designated as $1b_1, 1b_2, \dots, 1b_d$ with the bit position of each bit of the d sets of lower order bits designated from the most significant of the lower order bits to the least significant of the lower order bits by a value of i ranging, correspondingly, from n-1 to 0. The pruned tetrahedral interpolator includes a first multiplexer having a first multiplexer output and a first control input. The first multiplexer is configured for receiving the 2^d of the interpolation data values. The first control input is configured for receiving a first value determined from the d sets of lower order bits. The pruned tetrahedral interpolator further includes a first adder having a first input, a second input, and an output. The first input is coupled to the first multiplexer output and the second input is configured to receive one of the 2^d of the interpolation data values.
- 40 [0016] A method of pruned tetrahedral interpolation uses interpolation data values for selection using input data values each having d components. The d components are represented by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits with each of the d sets of lower order bits having n bits. The d sets of lower order bits are designated as $1b_1, 1b_2, \dots, 1b_d$ with the bit position of each bit of the d sets of lower order bits designated from the most significant of the lower order bits to the least significant of the lower order bits by a value of i ranging, correspondingly, from n-1 to 0. The method of pruned tetrahedral interpolation includes a step of computing a first set of $2^n - 2$ values using bitwise AND operations and bitwise OR operations operating upon $v[i]$. Where $v[i]$ is equal to $2^{d-1} \times 1b_1[i] + 2^{d-2} \times b_2[i] + \dots + 2^{d-d} \times 1b_d[i]$ for the value of i ranging from (n-1) to 0. The method of pruned tetrahedral interpolation further includes a step of selecting at least the minimum of 2^n and 2^d of the interpolation data values using ones of the first set of $2^n - 2$ values, $v[i]$ for one of the values of i equal n-1, and the d sets of higher order bits. The method of pruned tetrahedral interpolation further includes a step of adding a second set of the interpolation data values formed from the interpolation data values from the step of selecting to generate a sum.
- 45 [0017] A pruned tetrahedral interpolator for interpolating between interpolation data values uses input data values each having d components. The d components are represented by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits with each of the d sets of lower order bits having n bits. The d sets of lower order bits are designated as $1b_1, 1b_2, \dots, 1b_d$ with the bit position of each bit of the d sets of lower order bits designated from the most significant of the lower order bits to the least significant of the lower order bits by a value of i ranging, correspondingly, from n-1 to 0. The pruned tetrahedral interpolator includes a first multiplexer having a first multiplexer output and a first control input. The first multiplexer is configured for receiving the 2^d of the interpolation data values. The first control input is configured for receiving a first value determined from the d sets of lower order bits. The pruned tetrahedral interpolator further includes a first adder having a first input, a second input, and an output. The first input is coupled to the first multiplexer output and the second input is configured to receive one of the 2^d of the interpolation data values.
- 50 [0018] A method of pruned tetrahedral interpolation uses interpolation data values for selection using input data values each having d components. The d components are represented by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits with each of the d sets of lower order bits having n bits. The d sets of lower order bits are designated as $1b_1, 1b_2, \dots, 1b_d$ with the bit position of each bit of the d sets of lower order bits designated from the most significant of the lower order bits to the least significant of the lower order bits by a value of i ranging, correspondingly, from n-1 to 0. The method of pruned tetrahedral interpolation includes a step of computing a first set of $2^n - 2$ values using bitwise AND operations and bitwise OR operations operating upon $v[i]$. Where $v[i]$ is equal to $2^{d-1} \times 1b_1[i] + 2^{d-2} \times b_2[i] + \dots + 2^{d-d} \times 1b_d[i]$ for the value of i ranging from (n-1) to 0. The method of pruned tetrahedral interpolation further includes a step of selecting at least the minimum of 2^n and 2^d of the interpolation data values using ones of the first set of $2^n - 2$ values, $v[i]$ for one of the values of i equal n-1, and the d sets of higher order bits. The method of pruned tetrahedral interpolation further includes a step of adding a second set of the interpolation data values formed from the interpolation data values from the step of selecting to generate a sum.
- 55 [0019] A pruned tetrahedral interpolator for interpolating between interpolation data values uses input data values each having d components. The d components are represented by d sets of bits each partitioned to form d sets of higher order bits and d sets of lower order bits with each of the d sets of lower order bits having n bits. The d sets of lower order bits are designated as $1b_1, 1b_2, \dots, 1b_d$ with the bit position of each bit of the d sets of lower order bits designated from the most significant of the lower order bits to the least significant of the lower order bits by a value of i ranging, correspondingly, from n-1 to 0. The pruned tetrahedral interpolator includes a first multiplexer having a first multiplexer output and a first control input. The first multiplexer is configured for receiving the 2^d of the interpolation data values. The first control input is configured for receiving a first value determined from the d sets of lower order bits. The pruned tetrahedral interpolator further includes a first adder having a first input, a second input, and an output. The first input is coupled to the first multiplexer output and the second input is configured to receive one of the 2^d of the interpolation data values.

each having d components to generate output data values. The d components are represented by d sets of bits partitioned to form d sets of lower order bits with each of the d sets of lower order bits having n of the bits. The pruned tetrahedral interpolator includes a first set of $2^n - 1$ multiplexers each configured for receiving one of a set of control inputs and having a multiplexer output. Each of the multiplexers of the first set for selecting from the interpolation data values is responsive to the one of the set of control inputs. The pruned tetrahedral interpolator further includes a means for adding configured for receiving the multiplexer output of the set of multiplexers.

DESCRIPTION OF THE DRAWINGS

10 [0015] A more thorough understanding of the invention may be had from the consideration of the following detailed description taken in conjunction with the accompanying drawings in which:

Figure 1a is a representation of output color space values used for interpolation in a cubic lattice. The vertices of the each of the cubes forming the cubic lattice represent values of the output color space.

15 Figure 1b is a graphical representation of a color space conversion process from a color expressed in a cylindrical coordinate to a color expressed in a rectangular coordinate.

Figures 2a through 2d are a graphical representation of the selection of a single sub-cube using the corresponding bits from the lower order bits of the input color space value.

20 Figures 3a through 3h show the eight possible sub-cubes which can be selected from a cube using the corresponding bits from the lower order bits of the input color space value.

Figure 4 shows the numbering of the vertices of the cube for the purposes of selecting the sub-cube containing the result of the interpolation using the corresponding bits of the lower order bits.

Figure 5 is a graphical representation of the radial sub-cube generation process.

25 Figure 6a through 6e shows a graphical representation of multiple iterations of the cube subdivision process used in radial interpolation.

Figure 7 shows a hardware implementation of a radial interpolator.

Figure 8 is a high level flow diagram of a generalized method for performing radial interpolation.

Figure 9 is a diagrammatic representation of the computations required for generation of the sub-cubes in pruned radial interpolation.

30 Figure 10 shows a hardware implementation of pruned radial interpolation.

Figure 11 is a high level flow diagram of a generalized method for performing pruned radial interpolation.

Figure 12 is a high level flow diagram of a method implemented in software for performing pruned radial interpolation.

35 Figure 13 is a representation of the outer bounds of a CMY and a RGB color space showing the colors at the outer bounds of the color spaces.

Figure 14 is a graphical representation of the generation of a sub-cube from two tetrahedrons.

Figure 15 is a diagrammatic representation of the computations required for generation of the sub-cubes in pruned tetrahedral interpolation.

Figure 16 is a high level flow diagram of a generalized method for performing tetrahedral interpolation.

40 Figure 17 shows a hardware implementation of tetrahedral interpolation.

Figure 18 shows a hardware implementation of pruned tetrahedral interpolation.

Figure 19 is a high level flow diagram of a generalized method for performing pruned tetrahedral interpolation.

Figure 20 is a high level flow diagram of a method implemented in software for performing pruned tetrahedral interpolation.

45 Figure 21 is a diagrammatic representation of a common radial interpolation and pruned tetrahedral interpolation implementation.

Figure 22 shows a hardware implementation of common pruned radial and pruned tetrahedral interpolation.

Figure 23a through 23e shows a graphical representation of a non-symmetric interpolation process.

50 Figure 24 shows a graphical representation of the generation of a sub-cube from a cube using non-symmetric sub-cube generation.

Figure 25 is a diagrammatic representation of the non-symmetric radial interpolation process.

Figure 26 is a high level flow diagram of a method for performing non-symmetric radial interpolation.

Figure 27 is a high level flow diagram of a method implemented in software for performing non-symmetric pruned radial interpolation.

55 Figure 28 shows a hardware implementation of non-symmetric pruned radial interpolation.

Figure 29 is a high level flow diagram of a generalized method for performing non-symmetric radial interpolation.

Figure 30 shows a hardware implementation of non-symmetric radial interpolation.

Figure 31 is a diagrammatic representation of the non-symmetric pruned tetrahedral interpolation process.

Figure 32 shows a high level flow diagram of a method for performing non-symmetric pruned tetrahedral interpolation.

Figure 33 is a high level flow diagram of a method implemented in software to perform non-symmetric pruned tetrahedral interpolation.

5 Figure 34 shows a hardware implementation of non-symmetric pruned tetrahedral interpolation.

Figure 35 is a high level flow diagram of a generalized method for implementing non-symmetric tetrahedral interpolation.

Figure 36 shows a hardware implementation of a non-symmetric tetrahedral interpolator.

10 Figure 37 shows a hardware implementation of a common non-symmetric pruned radial and non-symmetric pruned tetrahedral interpolator.

Figure 38 includes a C code listing of a method for implementing pruned radial interpolation in software.

Figure 39 includes a VHDL listing used for generating a hardware implementation of pruned radial interpolation.

Figure 40 includes a C code listing of a method for implementing pruned tetrahedral interpolation in software.

15 Figure 41 includes a VHDL listing used for generating a hardware implementation of pruned tetrahedral interpolation.

Figure 42 includes a VHDL listing used for generating a hardware implementation of common pruned radial interpolation and pruned tetrahedral interpolation.

Figure 43 includes a C code listing of a method for implementing non-symmetric pruned radial interpolation in software.

20 Figure 44 includes a VHDL listing used for generating a hardware implementation of non-symmetric pruned radial interpolation.

Figure 45 includes a C code listing of a method for implementing non-symmetric pruned tetrahedral interpolation in software.

25 Figure 46 includes a VHDL listing used for generating a hardware implementation of non-symmetric pruned tetrahedral interpolation.

Figure 47 includes a VHDL listing used for generating a hardware implementation of common non-symmetric pruned radial and non-symmetric pruned tetrahedral interpolation.

DETAILED DESCRIPTION OF THE PREFERRED EMBODIMENT

30 [0016] The present invention is not limited to the specific exemplary embodiments illustrated herein. In addition, although several embodiments of sub-cube interpolation will be discussed in the context of a color laser printer, one of ordinary skill in the art will recognize after understanding this specification that the disclosed embodiments of sub-cube interpolation have applicability in any interpolative data transformation between spaces. For example, the interpolations required for the rendering of three dimensional graphics could advantageously use the disclosed interpolation techniques.

35 [0017] Sub-cube interpolation using tetrahedral interpolation to generate each of the vertices of the successive sub-cubes is taught in the co-pending application having USPTO serial number 08/504,406, the disclosure of which is incorporated by reference herein. However, the method for generation of the sub-cube vertex values disclosed in this co-pending application requires a large number of computations. A need exists for a method of generating the sub-cube vertex values which is more computationally efficient.

40 [0018] Shown in Figure 1a is a cubic lattice 1. The cubic lattice is formed of a multiplicity of cubes with the vertices of the cubes representing values in the output color space. The input color space values are each partitioned into an upper portion and a lower portion. The upper portion of each of the input color space values serves as an index to address the vertex values of the cubic lattice 1 used for interpolation. The lower portion of each of the input color space values is used to interpolate between the output color space values accessed using the upper portion of the input color space value. Each of the dimensions of the cubic lattice 1 correspond to one of the components of the input color space value. The values associated with the vertices of cubic lattice 1 are used to generate output color space values.

45 [0019] Each of the output color space values has multiple components corresponding to the dimensions of the output color space. Conversion is done from the input color space values to components of the output color space values. Conversion to each output color space value component uses a distinct set of vertex values. For the case in which there are three components to each of the output color space values, there are three sets of vertex values used for the color space conversion. For this case, it would be possible to regard each vertex value as formed of three values with each of the three values selected from one of the three sets. Viewing the vertex values in this way, the conversion to each of the components of the output color space values would be performed in parallel. It is also possible to perform the conversion to each of the output color space components serially. Done in this manner, the conversion can be viewed as using three separate cubic lattices, one corresponding to each set of vertex values.

55 [0020] Shown in Figure 1b is a general graphical representation of the interpolation process. Consider, for example

the conversion of an input color space value (a, b, c) 10, representing a color in a cylindrical color space, to an output color space value (x, y, z) 11 representing that same color in a Cartesian color space. In this example, each of a, b, and c are represented by eight bits. Each of the three groups of eight bits can be partitioned; for example, into four upper bits 10a (represented by a_u , b_u , and c_u) and four lower bits 10b (represented by a_1 , b_1 , and c_1). The three groups of four upper bits 10a are used as an index into the cubic lattice 1 to retrieve the eight values corresponding to the vertices of a cube within the cubic lattice 1 that will be used as interpolation data values. The three groups of four lower bits 10b are then used to interpolate between the eight interpolation data values corresponding to the vertices of the cubic lattice 1 to generate a component of the the output color space value 11.

5 [0021] One of ordinary skill in the art will recognize that other partitions of the bits of the input color space value 10 are possible. The particular partition of the bits will depend upon such things as the size of the memory available to store the values of the output color space used for the interpolation ad the amount of change in the output color space value that occurs between vertices of cubic lattice 1. A tradeoff exists between the accuracy of the interpolation and the size of the memory used to store the output color space values used as the interpolation data values. If the characteristics of the output color space are such that it changes relatively linearly throughout the color space, then fewer vertices 15 in cubic lattice 1 are necessary to deliver an acceptable level of interpolation accuracy.

[0022] The index formed by a_u , b_u , and c_u serves as an entry point into the cubic lattice. The index addresses one vertex of the eight vertices of the cube used as the interpolation data values. Each of the vertices of the cube corresponds to a value used for interpolating to generate one component of the output color space value 11. The eight associated vertices of a cube in the cubic lattice 1 have the following relative addresses:

20

- (a_u , b_u , c_u)
- (a_u , +1, b_u , c_u)
- (a_u , b_u , +1, c_u)
- (a_u , b_u , c_u + 1)
- 25 (a_u , +1, b_u + 1, c_u)
- (a_u , +1, b_u , c_u + 1)
- (a_u , b_u + 1, c_u + 1)
- (a_u , +1, b_u + 1, c_u + 1)

30 [0023] The cube subdivision interpolation method disclosed in the United States patent application having USPTO serial number 08/504,406, performs a interpolation by generating a sub-cube using the values associated with the vertices of the previously generated sub-cube. The initial cube formed by the vertex values associated with the three groups of upper order bits (a_u , b_u , c_u) 10a is used to generate the first sub-cube. This initial cube can be divided into eight sub-cubes. The three groups of lower order bits 10b (a_1 , b_1 , c_1) are used to select one of the eight possible sub-cubes formed for the next iteration of sub-cube division. These three groups of lower order bits 10b identify in which of the eight possible sub-cubes the result of the interpolation will be located. When the sub-cube which contains the result of the interpolation is identified, this sub-cube is used to generate the next sub-cube which contains the result of the interpolation. This process is successively repeated until the last sub-cube containing the result of the interpolation is generated. One of the values associated with a vertex of this last sub-cube generated is used as the result of the interpolation.

40 [0024] Figure 2a through 2d graphically represent the selection of a sub-cube using the three groups of lower order bits 10b (a_1 , b_1 , c_1). For purposes of explaining the sub-cube selection, consider the case in which the lower order bits 10b for each component of the input color space value consists of four bits. Shown in each of Figures 2a through 2d are the axes corresponding to the a, b, and c components of the input color space value. Each of these axes corresponds 45 to a dimension of the input color space. Sub-cubes are designated using one corresponding bit (corresponding in the sense that they are coefficients of the same power of 2) from the lower order bits 10b of each component of the input color space value. Each bit position of the lower order bits 10b of each component can be viewed as dividing the cube in half along the dimension corresponding to the component. The value of the bit for each component determines which half of the cube is selected in the corresponding dimension, for the purpose of determining in which sub-cube the result 50 of the interpolation is located. The selected sub-cube will be the volume defined by the intersection of the cube halves selected by the corresponding bits of each component of the lower order bits 10b of the input color space value. If the bit of the lower order bits 10b for the component is a "1", a corner of the selected cube half is displaced one half the length along the corresponding axis from the origin of the cube. If the bit of the component is a "0", the corner of the selected cube half includes the origin of the cube. Shown in Figures 3a through 3h are the eight possible sub-cubes 55 defined by the common intersection of the cube halves designated by the corresponding bit from each of a_1 , b_1 , and c_1 . [0025] By numbering the vertices of the cubes in a manner that is consistent with the assignment of the groups of lower order bits 10b of the components of the input color space value to the axes, the vertex of the cube used to generate the sub-cube associated with a vertex of the sub-cube to be generated, is designated by the binary value formed

EP 0 923 049 A2

by combining the corresponding bits from each of the pups of lower order bits 10b of the components. Shown in Figure 4 is a cube with the axes labeled and with the vertices numbered. The cube used to generate a sub-cube and the generated sub-cube share a vertex. With this assignment of vertex numbers, the number of the vertex of the cube used to generate the sub-cube which is included within the sub-cube generated is the binary value formed from the corresponding bits of a_1 , b_1 , and c_1 for a given bit position.

[0026] An example will be explained to illustrate the interpolation using the sub-cube generation. Assume that the following values are used for a_1 , b_1 , and c_1 :

10

Table 1

$$a_1=1010$$

$$b_1=1110$$

$$c_1=0011$$

15

With these values assigned to a_1 , b_1 , and c_1 , the vertex number 6 (computed by selecting the most significant bit from each of a_1 , b_1 , and c_1 and concatenating these into a binary value) of the cube used to generate the first sub-cube is also a vertex of the first sub-cube. The vertex of the first sub-cube included within the second sub-cube generated is vertex number 2 (computed by selecting the second most significant bit from each of a_1 , b_1 , and c_1 and concatenating these into a binary value). The vertex of the second sub-cube included with the third sub-cube generated is vertex number 7 (computed by selecting the third most significant bit from each of a_1 , b_1 , and c_1 and concatenating these into a binary value). The vertex of the third sub-cube included with the fourth sub-cube generated is vertex number 1 (computed by selecting the fourth most significant bit from each of a_1 , b_1 , and c_1 and concatenating these into a binary value). The vertex values of the first sub-cube are generated with the vertex values accessed using the upper order bits (a_u , b_u , c_u) 10a of the components of the input color space value. The vertex values of the second sub-cube are generated using the vertex values generated for the first sub-cube. The vertex values of the third sub-cube are generated using the vertex values generated for the second sub-cube. Finally, the vertex values of the fourth sub-cube are generated using the vertex values generated for the third sub-cube. In sub-cube interpolation, the value associated with the vertex numbered 0 of the final sub-cube generated is the value used as the result of the interpolation. This result is one component of the output color space value. This sub-cube generation procedure could be applied with an arbitrary number of bits used to specify each component of the lower order bits 10b of the input color space value.

[0027] A variety of methods have been previously employed for the generation of the sub-cube values. These methods include tetrahedral, pyramid, PRISM, and Trilinear. Radial sub-cube generation is a new method of sub-cube generation which achieves a substantial reduction in the computational complexity required to generate the sub-cubes. It should be recognized that each interpolation method can generate different results because the interpolation process is an approximation of the color space conversion. Depending upon the location in the color space where the conversion is performed and the preferred characteristics of the result, one method may yield more desirable results than another.

[0028] Shown in Figure 5 is a graphical representation of the radial sub-cube generation method. Clearly explaining the radial sub-cube generation process requires some notational definition. As was previously the case, a_1 , b_1 , and c_1 designate the lower order bits 10b of the respective a, b, and c components of input color space value. The value of the variable i will be used to designate the bit position within the lower order bits 10b (a_1 , b_1 , c_1), as shown below, for the case in which four bits are used to designate each component of the lower order bits 10b. The maximum value of i (a value of 3) corresponds to the most significant bit position of the lower order bits 10b. The minimum value of i (a value of 0) corresponds to the least significant bit position of the lower order bits 10b. As one of ordinary skill in the art will recognize, this notation is easily adapted for a different number of bits used for each component of the lower order bits 10b. For n bits used to represent the lower order bits 10b, the value of i ranges from $n-1$ to 0.

50

Table 2

55

i:	3	2	1	0
a_1 :	1	0	1	0
b_1 :	1	1	1	0
c_1 :	0	0	1	1

[0029] Using this notation, the value of i indirectly indicates the iteration of the sub-cube generation. A value of i equal to 3, corresponds to generation of the first sub-cube. This first sub-cube includes vertex number 6 from the cube formed by accessing the values of cubic lattice 1 using the upper order bits 10a. For the value of i equal to 0, the fourth sub-cube is generated. This fourth sub-cube includes vertex number 1 from the third sub-cube generated. To determine the vertex number of the cube used to generate the sub-cube which is included within the sub-cube, the following equation is used:

$$v(i) = 4a_1(i) + 2b_1(i) + c_1(i) \quad \text{eqn. 1a}$$

In equation 1a, $v(i)$ represents the vertex number of the cube included within the generated sub-cube. Each of the $a_1(i)$, $b_1(i)$, and $c_1(i)$ represents the binary value associated with the "ith" position in the respective component of the lower order bits 10b. For each value of i , equation 1a yields the correct number of the vertex of the cube used to generate the sub-cube which will be included within the desired sub-cube. The values which i may assume include the integers from $n - 1$ to 0 inclusive, where n is the number of bits used to specify each of the components of the lower order bits 10b of the input color space value. The value of the vertex having the number $v(i)$ is designated by $P[v(i)]$. Equation 1a can be generalized for input color space value 10 formed from d components. Given below is a generalized expression for $v[i]$:

$$v[i] = 2^{d-1} \times 1b_1[i] + 2^{d-2} \times 1b_2[i] + 2^{d-3} \times 1b_3[i] + \dots + 2^{d-d} \times 1b_d[i] \quad \text{eqn. 1b}$$

In equation 1b, each of the "1b" represent the lower order bits 10b one of the d components of the input color space value 10. As in equation 1a, the values of i include the integers from $n-1$ to 0 inclusive.

[0030] The values associated with the eight vertices of the sub-cube are generated from the cube as shown in table 3. Those vertex values designated by $P[\text{sub-cube vertex number}]$ represent sub-cube vertex values and those vertex values designated by $P[\text{cube vertex number}]$ represent the vertex values of the cube used to generate the sub-cube vertex values. A given value of a sub-cube vertex is generated by averaging the corresponding vertex value of the cube from which the sub-cube is generated with the value of the vertex of the cube used to generate the sub-cube included within the sub-cube (this is the value designated by $P[v(i)]$).

30

Table 3

$$\begin{aligned} P[7] &= (P[7] + P[v(i)]) / 2 \\ P[6] &= (P[6] + P[v(i)]) / 2 \\ P[5] &= (P[5] + P[v(i)]) / 2 \\ P[4] &= (P[4] + P[v(i)]) / 2 \\ P[3] &= (P[3] + P[v(i)]) / 2 \\ P[2] &= (P[2] + P[v(i)]) / 2 \\ P[1] &= (P[1] + P[v(i)]) / 2 \\ P[0] &= (P[0] + P[v(i)]) / 2 \end{aligned}$$

45

[0031] Shown in Figure 6a through Figure 6e is a graphical representation of multiple iterations of the sub-cube division process used in radial interpolation. The values used for the components of the lower order bits 10b of the input color space value in the example of Figure 6 are the same as those shown in Table 2. The values of the vertices of the cube used for generation of the first sub-cube (these values are accessed using the higher order bits 10a of the input color space value) are loaded from a color table stored in memory. After the final iteration of sub-cube division, vertex number 0 of the final sub-cube is used as the result of the interpolation process. To prevent the accumulation of rounding errors during the actual computation of the sub-cube vertex values, the required division by 2 for each iteration of sub-cube generation is performed only on the values of the vertices of the final sub-cube generated. When the division is performed in this manner, the divisor used is 2^n , where n is the number of bits assigned to each component of the lower order bits 10b. Division by 2^n can be performed easily by performing a right shift operation. For the case in which $n=4$, this divisor is 16. With the division operation not performed until after the generation of the final sub-cube vertex values, the sub-cube generation process reduces to a series of additions of selected vertex values of the generated sub-cubes.

[0032] Shown in Figure 7 is a hardware implementation of a radial interpolator 100. In addition, Figure 7 illustrates the progression of the radial interpolation through radial interpolator 100 using values for $v[i]$ corresponding to lower order bits 10b of table 2. Each of the values ($P[0]$ through $P[7]$) associated with the eight vertices selected using upper order bits 10a is coupled to a multiplexer input of first multiplexer 101. The value of $v[i]$, for i equal to 3, is coupled to the control input of first multiplexer 101. The value of $v[i]$, for i equal to 3, is used to select the value associated with the vertex of the cube selected using upper order bits 10a that will be included in the first sub-cube generated. The output of first multiplexer 101 is coupled to a first input of each adder of a first set of adders 102 composed of eight adders. The second input of each adder of the first set of adders 102 is coupled to one of the values selected using upper order bits 10a. First multiplexer 101 and first set of adders 102 form a first stage of radial interpolator 100. It can be seen, that with this configuration of multiplexer 101 and first set of adders 102, the averaging operations of table 3 (without the division by two, which, as previously mentioned, is delayed until all the iterations of radial interpolation are completed) for a single iteration of radial interpolation are completed. A second, third, and fourth stage of radial interpolator 100 are formed from, respectively, a second multiplexer 103 and second set of adders 104, a third multiplexer 105 and third set of adders 106, and a fourth multiplexer 107 and fourth set of adders 108. The control inputs of the second 103, third 105, and fourth 107 multiplexer inputs are coupled to, respectively, $v[i=2]$, $v[i=1]$, and $v[i=0]$. The second, third, and fourth stages of radial interpolator perform successive iterations of radial interpolation with each iteration using the relationships of table 3 (again delaying the division by two until completion of all iterations).

[0033] Shown in Figure 8 is high level flow diagram of a method for performing a single iteration of radial interpolation. First, 2^d of interpolation data values are selected 200 using upper order bits 10a. For the case in which the radial interpolation is used for color space conversion, the interpolation data values correspond to vertex values of the selected cube. In addition, for color space conversion d is typically equal to 3, the number of components of the input color space value 10. After vertex values are selected 200, the vertex number of the vertex value required for that iteration is computed 201. For d equals 3, equation 1a is used to compute 201 the required vertex number. Depending on the iteration of radial interpolation, the value of i used to compute $v[i]$ can range from $i=3$ to $i=0$. After computation 201 of the vertex number, one of the 2^d of interpolation data values is selected 202 using the computed vertex number. Finally a set of 2^d averages is computed 203 according to the relationships of table 3. To avoid rounding errors the required divisions by 2 for averaging are delayed until all iterations of interpolation are performed.

[0034] It was recognized that the number of computations required to perform radial interpolation as shown in Figure 7 could be substantially reduced. Examination of the radial interpolation process of Figure 7 reveals that determination of the interpolation result does not require the use of all eight of the vertex values accessed by the upper order bits 10a, nor does it require the use of all the adders shown in Figure 7. Shown in Figure 9 is a diagrammatic representation of pruned radial interpolation. Figure 9 can be understood by working backward from the interpolation result shown in Figure 7 to determine the values of the vertices (accessed using the upper order bits 10a) that are required to generate the result. As previously mentioned, the value of vertex number 0 ($P[0]$) of the final sub-cube generated is used as the result of the interpolation. Using the equations listed in table 3 (without the division by 2), the vertices of the sub-cube immediately previous to the final sub-cube that are used to compute the value of vertex number 0 of the final sub-cube can be determined. Similarly, the equations listed in table 3 can be used to determine the vertices of the sub-cube two previous to the final sub-cube that are necessary to compute the needed vertices of the sub-cube immediately previous to the final sub-cube.

[0035] This method for determining the vertices of each of the sub-cubes necessary to compute $P[0]$ of the final sub-cube is performed with each value of i from 0 to 3 for the case in which $n=4$. If this is done, the result shows that the values used to compute $P[0]$ of the final sub-cube consist of only some of the values corresponding to the vertices of the cube accessed by the higher order bits 10a of the input color space value. As a result, only 10 of the 32 adders of Figure 7 are used for computing a interpolation result from a given input color space value 10. For the values of the lower order bits 10b shown in table 2, the vertex numbers of the cube accessed by the higher order bits 10a to which the values used to compute $P[0]$ of the final sub-cube correspond are: 0, 6, 2, 7, and 1. The values of the vertices of the cube accessed by the higher order bits 10b corresponding to this are: $P[0]$, $P[6]$, $P[2]$, $P[7]$, and $P[1]$. In general, the values of the vertices of the cube accessed using the higher order bits 10a that are used to compute $P[0]$ of the final sub-cube are $P[0]$, $P[v(i=3)]$, $P[v(i=2)]$, $P[v(i=1)]$, $P[v(i=0)]$. The general expression which can be derived for $n=4$ is:

$$P[0]_{\text{Final Sub-cube}} = \{(8 \times P[v(i=3)]) + (4 \times P[v(i=2)]) + (2 \times P[v(i=1)]) + P[v(i=0)] + P[0]\} \div 16 \quad \text{eqn. 2}$$

[0036] Equation 2 is an expression for computing a result using pruned radial interpolation with $n=4$. A generalized expression for the pruned radial interpolation is:

$$P[0]_{\text{Final Sub-cube}} = \left\{ \sum_{i=0}^{n-1} (2^i \times P[v(i)]) + P[0] \right\} \div 2^n \quad \text{eqn. 3}$$

Equation 3 can be used to generalize the computation of the pruned radial interpolation result. It should be noted that in equation 2 and the generalized expression in equation 3, the value associated with vertex number 0 of the cube selected using higher order bits 10a is always used. Had the value of a vertex number other than vertex number 0 of the final sub-cube generated been used as the result of the interpolation, the value of that vertex number of the originally selected cube would be used in place of P[0].

- 5 [0037] The hardware functional blocks required to perform the pruned radial interpolation include adders and multiplexers. With D dimensions in the output color space and bits representing each group of lower order bits 10b of the input color space value, the requirements of the hardware implementation of the pruned radial interpolation can be computed as:

10 $\# \text{ of Adders} = D \times (n + 1)$ eqn. 4

$\# \text{ of Multiplexers} = D \times n$ eqn. 5

- 15 [0038] It should be noted that extra adder specified in equation 4 is used for the purpose of rounding. Additional operations which must be performed by the hardware include multiplication, division, and concatenation. The multiplication and division operations by a power of two can be performed by shifting bit positions. In hardware, this shifting is accomplished by connecting a line corresponding to a bit to a higher order position for multiplication or to a lower order position for division. In hardware, concatenation is accomplished by grouping lines, corresponding to bit positions, together.

- 20 Therefore, the multiplication, division, and concatenation operations can be performed without the necessity of adding additional hardware.

- [0039] To generate the gate level design necessary to implement the pruned radial interpolation in hardware, a commonly used hardware description language, such as VHDL, may be used. Included in Figure 39 is a listing of the VHDL code which can generate a hardware implementation of pruned radial interpolation.

- 25 [0040] Shown in Figure 10 is a hardware implementation of a pruned radial interpolator 300 for n equal to 4. It should be noted that the hardware implementation shown in Figure 10 can be used to generate a single component of the output color space value 11. This same hardware could be used repetitively for an additional (D-1) passes to generate the remaining D-1 components of the output color space value 11. Or, there could be an additional (D-1) replications of the hardware implementation shown in Figure 10 to generate each of D components simultaneously. Pruned radial interpolator 300 is a hardware implementation of equation 2 without the division by 16. The division by 16 could be accomplished by bit shifting the result of the additions. Selection of four of the five vertex values (P[6], P[2], P[7], and P[1]) used to compute the interpolation result requires four multiplexers 301-304.

- 30 [0041] The four required additions are accomplished using four adders 305-308. The fifth vertex value required for computation of the interpolation result, P[0], is hard wired into the inputs of one of the adders. It should be noted that because of the associative property of addition, the hardware implementation of Figure 10 may be implemented so that the additions are performed in a number of different orders. The order shown in Figure 10 minimizes propagation delay through the adders. Furthermore, other means for adding may be used. For example, a single adder that had a sufficient number of inputs could be used. Or, a microprocessor could be used to accomplish the additions.

- 35 [0042] The three multiplication operations 309-311 correspond to multiplication by the coefficients 8, 4, and 2 of the first three terms on the right side of equation 2. It should be noted that multiplication operations 309-311 are accomplished in hardware by routing of the lines corresponding to the bit positions on each of the respective multiplexer outputs. Therefore, these multiplications are implemented without additional hardware cost. It is possible to implement the multiplications through the routing of lines because all of the coefficients are powers of 2.

- 40 [0043] One of ordinary skill in the art will recognize that the hardware implementation shown in Figure 10 is adaptable for values of n greater than 4 or less than 4. Consider the hardware implementation of a pruned radial interpolator for n equal to 1. This hardware implementation of pruned radial interpolation would be useful for an interpolation which performs a single iteration of cube subdivision and then selects one of the vertex values of the generated sub-cube as the interpolation result. This hardware implementation of pruned radial interpolation requires only a single multiplexer and a single adder (of course the rounding at the end requires an additional adder but this additional adder is not shown in

- 45 Figure 10).
- [0044] Shown in Figure 11 is a high level flow diagram of a generalized method of pruned radial interpolation for input color space values 10 having d components with each set of lower order bits 10b having n bits. First, n values are computed 400 using equation 1b. Next, n+1 interpolation data values (which correspond to vertex values in a color space conversion) are selected 401 using the computed n values and higher order bits 10a. Finally, the interpolation result is computed 402 by multiplying and adding the selected n+1 interpolation data values according to equation 3.

- 50 [0045] A software implementation of the pruned radial interpolation is computationally very efficient. With d input dimensions, D output dimensions and 2^n values between vertices of cubic lattice 1, the number of computations required to generate an interpolation result can be computed as:

$$\# \text{ of ALU operations} = 2 \times n \times (d + D - 1) + D$$

eqn. 6

The number of memory accesses required to generate the interpolation result can be computed as:

5

$$\# \text{ of memory accesses} = D \times (n + 1)$$

eqn. 7

It should be noted that, unlike many other interpolation methods, both the number of ALU operations and the number of memory accesses are linear in D, d, and, n which results in the relative computational efficiency of pruned radial interpolation.

10 [0046] Shown in Figure 12 is a high level flow diagram of a method implemented in software to perform the pruned radial interpolation. First, a determination 500 is made if any one of the components of the input color space value (a, b, c) 10 corresponds to a location on an outer boundary of the cubic lattice 1. This is the case if any one or more of the components of the input color space value has a value of FF hexadecimal. If this is the case, then, for purposes of generating the index into the cubic lattice 1 to retrieve the necessary vertex values, the components of the input color space value 10 which have a value of FF hexadecimal are assigned 501 a value of 100 hexadecimal.

15 [0047] Assignment of a value of 100 hexadecimal to those input color space values of FF hexadecimal is done to address a special case in the interpolation. To illustrate this special case, consider the representation of the input color space values 10 by eight bits for each component, with each component partitioned into four upper order bits and four lower order bits. With this partitioning, the higher order bits can form the index values 00, 10, 20, 30,...F0 hexadecimal 20 for each component. The four lower order bits for each component will be used to interpolate between the output color space values 11 accessed using the index values. The difference between the pair of output color space values 11 accessed using successive index values from 00 hexadecimal to F0 hexadecimal is spanned in 16 equal increments. With each successive increment, the associated value is increased 1/16 of the difference between the accessed pair of output color space values 11, when going from the lower output color space value to the higher output color space 25 value. For example, after 5 increments, the associated value is 5/16 of the difference between the accessed pair of output color space values 11. Using the four lower order bits, the value associated with the corresponding number of increments is added to the output color space value 11 selected using the higher order bits to generate the interpolation result. However a problem arises between index values F0 and FF (index value 100 does not exist in the table) for each component of the input color space value 10. Between F0 and FF there are only 15 increments and the output color 30 space value 11 accessed by FF corresponds to an outer boundary of the output color space. However, the interpolation process is designed to operate on 16 increments between the output color space values 11 accessed using the index values. To address this problem, the output color space values 11 corresponding to the index value FF are mapped to a location having an address of 100 hexadecimal. This mapping effectively distributes the difference in the output color space values 11 corresponding to index values F0 and FF hexadecimal over 16 increments instead of 15. Because of 35 this, there will be slight errors resulting from the interpolation between index values F0 and FF. Although not shown in the hardware block diagrams, the handling of this special case in the interpolation is performed in the hardware implementations of the various interpolator embodiments.

35 [0048] After any necessary reassignment of input color space value 10, the indices used to access the values corresponding to the required vertices of the selected cube in cubic lattice 1 are computed 502. Finally, the values for each 40 component of the output color space value (x, y, z) 11 are computed 503. Provided in Figure 38 of this specification is the code of an implementation in C, for n=4, of the high level method of pruned radial interpolation shown in Figure 12. [0049] It should be recognized that a number of possible processor specific optimizations of the software for performing pruned radial interpolation can be performed. For example, by combining all the components of each output color 45 space values 11 into a single word, the number of memory accesses required to perform the conversion to the output color space value 11 can be reduced. Another possible optimization exploits the ability of the ALU to perform 32 bit operations. By assigning bits 0-7 of an ALU word to handle the computation of the y component of the output color space value and bits 16-23 to handle the computation of the x component of the output color space value, a single sequence of shifts and adds can be used to generate the x and y components in parallel. It is also possible to implement pruned radial interpolation in hardware. The computational efficiencies which existed in the software implementation of 50 pruned radial interpolation are present in the hardware implementation as reduced hardware requirements.

[0050] Tetrahedral interpolation partitions the cube accessed by the higher order bits 10a of the color space input value into a number of tetrahedrons used for generation of the sub-cube containing the result of the interpolation. The resulting sub-cube is then partitioned into tetrahedrons. Two of these tetrahedrons are then used to generate yet another sub-cube containing the result of the interpolation. The successive division of generated sub-cubes into tetrahedrons is performed n times, where n is the number of bits used to represent each of the components of the lower order bits 10b of the input color space value 10.

55 [0051] Shown in Figure 13 is representation of the outer bounds of a CMY or a RGB color space. As can be seen from Figure 13, the vertices of the cube 600 formed by the outer bounds of these color spaces include values corre-

sponding to the constituent colors of each of the color spaces. A characteristic of the CMY and RGB color spaces is that the diagonal connected between the white 601 and black 602 vertices of the color space corresponds to the luminance axis. Points along the luminance axis have values which correspond to various shades of gray. As previously mentioned, the higher order bits 10a of the input color space value 10 are used to access eight associated values forming a cube located within cube 600. Analogous to the cube 600 representing the CMY or RGB color space, each of the selected cubes can be regarded as a kind of miniature color space, with the values corresponding to each of the eight vertices having colors which are weighted toward the colors of the corresponding vertices of cube 600. For example, the vertex of the selected cube spatially corresponding to the yellow vertex 603 is the vertex having a value closest to the value for the color yellow within in the selected cube. The other seven vertices of the selected cube can be viewed similarly. The diagonal connecting vertex 0 and vertex 7 serves to define a constant chrominance line between the colors associated with the vertices of the selected cube.

[0052] Certain artifacts can arise from the reproduction of colors in the printing process. These artifacts are visually perceptible as colors which deviate from those specified by the color space value input to the printing process. The artifacts are particularly noticeable for input color space values located near the luminance axis. Input color space values near the luminance axis correspond to shades of gray with small amounts of color. Factors in the color reproduction process which may push the resulting color farther off the luminance axis than intended are easily perceived in a gray field. The artifacts appear as colors in fields which should include only various shades of gray along the luminance axis. The artifacts can arise from, among other things, the characteristics of the process used for printing (such as an electrophotographic or inkjet printing process) or characteristics of the colorants (such as toner or ink) used in the printing process. Variability in the parameters of the printing process result in the reproduction of colors off the luminance axis when the result should have been gray.

[0053] Tetrahedral interpolation, in some circumstances, reduces the degree to which these types of artifacts are perceptible. The reduction in print artifacts occurs because the value of one vertex of the sub-cube generated from the tetrahedron is computed using the values associated with vertex number 0 and vertex number 7. As previously mentioned, the diagonal formed between vertex number 0 and vertex number 7 defines a constant chrominance line for the selected cube. Computing a vertex of the sub-cube along this mid-point color boundary line produces a weighting in the interpolation which tends to reduce the rate of change in the output color space value as the input color space value 10 moves off the diagonal of the cube selected by the higher order bits 10a. This in turn tends to somewhat compensate for the variability in printing process parameters which produce non-gray output with gray input color space values.

[0054] Shown in Figure 14 is a graphical representation of the generation of a sub-cube 700 from a tetrahedron 701. Each value of a vertex of the tetrahedron 701 used to compute the value of a vertex of the sub-cube 700 is also a value of a vertex of the cube 702 from which the tetrahedron 701 was partitioned. Let $P[k]$ denote the value associated with vertex k of a cube. Let $P'[k]$ denote the value associated with vertex k of a sub-cube included within the cube having vertex k . It can be shown that the value $P'[k]$ is computed as:

$$P'[k] = [P[k \& v(i)] + P[k|v(i)]] \div 2 \quad \text{eqn. 8}$$

"&" represents the bitwise AND operation

"|" represents the bitwise OR operation

40 "k" represents the vertex number

$v(i) = 4 \times a(i) + 2 \times b(i) + c(i)$

i represents the bit position in the lower order bits 10b in the input color space value

[0055] The sub-cube generation that can be accomplished using equation 8 provides a new method of computing vertex values for a tetrahedral interpolation. By using equation 8, the indices used to access the values of the vertices used for computing the sub-cube vertex values can be computed. This provides an advantage over interpolation methods that require accessing of a look-up table in order to determine the indices used to access the vertex values. The use of a look-up table requires memory accesses. As a result, using a look-up table to generate the indices requires a significantly greater number of machine cycles than would be required using the processor to compute the indices. Therefore, using equation 8 to compute the indices used to access the vertex values provides a substantial speed advantage in tetrahedral interpolation over previous methods of performing tetrahedral interpolation. Furthermore, implementing equation 8 in hardware for computation of the multiplexer control inputs used to select the vertex values provides a simpler hardware implementation of tetrahedral interpolation.

[0056] As was the case for pruned radial interpolation, the value of vertex number 0 of the last sub-cube generated is the result in the tetrahedral interpolation. It was recognized that not all the values of the vertices of all of the sub-cubes generated were required to generate the result of the interpolation. This led to the development of a further improvement in tetrahedral interpolation referred to as pruned tetrahedral interpolation. Shown in Figure 15 is a dia-grammatic representation of pruned tetrahedral interpolation. In Figure 15, the prime indicator associated with the term

representing the value of each vertex indicates the level of sub-cube generation. For example, terms designated as $P'[]$ represent vertex values after the first cube subdivision iteration, terms designated as $P''[]$ represent vertex values after the second cube subdivision iteration. This method of designating vertex values applies for the generation of successive sub-cubes. Figure 15 represents the pruned tetrahedral interpolation using 4 bits for the lower order bits 10b of the input color space value 10. The terms shown in Figure 15 can be generated by starting with the end result of the interpolation $P'''[0]$, and determining, successively, using equation 8, the values of the vertices of the previous sub-cube required to generate the values of the vertices of the current sub-cube until the values required to generate the current sub-cube are obtained by accessing the values of the vertices in cubic lattice 1 using higher order bits 10a. As was the case for pruned radial interpolation, the divide by 2 operation is not performed until the value of vertex number 0 of the final sub-cube is obtained in order to prevent the accumulation of rounding errors.

[0057] Pruned tetrahedral interpolation provides a substantial computational savings over tetrahedral interpolation. With d input dimensions, D output dimensions, and n lower order bits 10b, the number of computations required to perform the pruned tetrahedral interpolation is computed as:

$$15 \quad \# \text{ of ALU Operations} = n \times [(2 \times d) - 1] + 2 \times \left[\sum_{i=1}^{n-1} (2^i - 1) + D \times (2^n + 1) \right] \quad \text{eqn. 9}$$

$$20 \quad \# \text{ of Memory Accesses} = D \times [\min(2^d, 2^n)] \quad \text{eqn. 10}$$

It should be noted that, for these equations, the computations vary linearly as a function of d and D for the number of ALU operations, exponentially as a function of d for the number of memory accesses, and exponentially as a function of n for the number of memory accesses. As shown in Figure 15, pruned tetrahedral interpolation is implemented so that 2^n memory references are required. However, it is possible to implement pruned tetrahedral interpolation so that the number of memory references required is a maximum of 2^d , which for d=3, is the number of vertices within a cube. This is done by recognizing that redundancy exists in the 2^n memory accesses. By using some of the accessed vertex values for multiple of the input values required in Figure 15 fewer memory accesses are required. Comparing equations 9 and 10 with equations 4 and 5 it can be seen that, with all other things equal, pruned tetrahedral interpolation is more computationally costly than radial interpolation.

[0058] Shown in Figure 16 is a high level flow diagram of a generalized method for performing tetrahedral interpolation. For this method, the input color space values 10 are formed of d components. Each of the d components is partitioned into a set of higher order bits 10a and lower order bits 10b. Each of the d sets of lower order bits 10b is formed of n bits. The d sets of lower order bits are each designated as 1b₁, 1b₂, 1b₃, ..., 1b_d. The bit position of each of the lower order bits is designated from the most significant bit to the least significant bit by a value of i ranging, correspondingly from n-1 to 0. First, a value is computed 800 according to $v[i] = 2^{d-1} \times 1b_1[i] + 2^{d-2} \times 1b_2[i] + 2^{d-3} \times 1b_3[i] + \dots + 2^{d-d} \times 1b_d[i]$ for i equal to n-1. Next, a set of AND values is computed 800 according to $v[i] \& k$, for the value of k ranging from 2^{d-1} to 0, where "&" represents the bitwise AND operation. Then, a set of OR values is computed 802 according to $v[i] | k$ for the value of k ranging from 2^{d-1} to 0, where "|" represents the bitwise OR operation. Next, 2^d pairs of the vertex values are selected 803 using the set of AND values and the set of OR values. Each of the pairs are selected using an AND value and an OR value computed for a corresponding value of k. Finally, a set of 2^d sums is computed 804 by summing each of the 2^d pairs of vertex values.

[0059] The method shown in Figure 16 is for a single iteration of tetrahedral interpolation. Performing successive iterations would require computing additional values of v[i], computing additional AND and OR values, selecting values from 2^d sums computed in the previous iteration using the additionally computed AND and OR values, and computing additional sets of 2^d sums. After the final iteration of tetrahedral interpolation, each of the final 2^d sums is divided by 2^n (not shown in Figure 16), where n is the number of iterations, and one of the resulting values is selected as the result of the interpolation. The division by 2^n is done after the final iteration, instead of dividing by two after each iteration, to prevent round-off error accumulation.

[0060] Shown in Figure 17 is a hardware implementation of a tetrahedral interpolator 900 for conversion of input color space values 10 to a component of output color space values 11. This same hardware could be used repetitively for an additional (D-1) passes to generate the remaining D-1 components of the output color space value 11. Or, there could be an additional (D-1) replications of part of the hardware implementation shown in Figure 17 to generate each of D components simultaneously. The hardware used for generating multiplexer control inputs could be used for each of the D replications. The tetrahedral interpolator 900 shown in Figure 17 corresponds to d=3 and n=4 for the input color space value 10. The hardware implementation shown in Figure 17 implements equation 8 for the generation of the sub-cube vertex values. The tetrahedral interpolator 900 of Figure 17 is formed from a first, second, third, and fourth stage 901-904. Each of the four stages 901-904 includes 2^3 adders, one of which is labeled as 905. Each of the four stages

901-904 further includes 2×2^3 multiplexers, one of which is labeled as 906, arranged as 2^3 pairs of multiplexers. Finally, each of the four stages 901-904 includes 2^3 bitwise OR blocks, one of which is labeled as 907, and 2^3 bitwise AND blocks, one of which is labeled as 908.

[0061] Each of stages 901-904 performs an iteration of interpolation. Some interpolation applications may require that 5 only a single iteration of interpolation be performed. For a single iteration of interpolation n=1. This corresponds to a hardware implementation of tetrahedral interpolator 900 using only first stage 901. An additional stage would be added for each additional iteration of interpolation required for the particular application.

[0062] Each of the inputs of the multiplexers in the first stage 901 are connected to the eight vertex values selected 10 using higher order bits 10a. The two outputs of each pair of multiplexers in first stage 901 are connected to the first and second inputs of the corresponding adder. The output of each of the adders of the first stage 901 is the vertex value of the first sub-cube. As previously mentioned, the division by two for each iteration of sub-cube generation is deferred until the last sub-cube is generated. The vertex values of the last sub-cube generated are divided by 2^n , where n is the number of bits in the lower order bits of the input color space value and n corresponds to the number of stages in the tetrahedral interpolator. The inputs of each multiplexer for the second, third, and fourth stages 902-904 are coupled to 15 the outputs of the adders of the previous stage.

[0063] The control input of one the multiplexers of each pair of multiplexers is connected to the output of a bitwise OR block. The control input of the other one of each pair of multiplexers is connected to the output of a bitwise AND block. The multiplexers used in the tetrahedral interpolator 900 have the capability to select one of eight, eight bit values using a three bit control input. The bitwise OR blocks and the bitwise AND blocks each perform, respectively, bit by bit OR 20 operations or AND operations on the values input to them. For this d=3 implementation of tetrahedral interpolator 900, each of the inputs to the bitwise OR blocks and bitwise AND is a 3 bit quantity. The output of each of the bitwise OR blocks and bitwise AND blocks to each of the multiplexers is a 3 bit quantity. The adders associated with each pair of multiplexers performs an addition of the selected eight bit values from each of the multiplexers.

[0064] As indicated by equation 8, the vertex number corresponding to the vertex value generated is connected to 25 one of the inputs for each corresponding bitwise OR block and bitwise AND block. Because these values are fixed they can be hardwired to the correct values. The other inputs for each corresponding pair of bitwise AND blocks and bitwise OR blocks in a stage are connected to the value of v[i] corresponding to the stage. For the first stage 901, the value is v[3]. For the second stage 902, the value is v[2]. For the third stage 903, the value is v[1]. For the fourth stage 904, the value is v[0].

[0065] Interpolation is performed by supplying the vertex values selected using higher order bits 10a to the multiplexer 30 inputs of the first stage 900 and supplying the appropriate v[i] values to the bitwise OR blocks and bitwise AND blocks of each stage. The tetrahedral interpolator 900 computes the vertex values for four iterations of sub-cube generation. The values P'''(7) through P'''(0) are the values of the vertices of the fourth sub-cube generated. In this embodiment, P'''(0) is selected, divided by 16, and used as the result of the interpolation. The division by 16 is implemented by shifting bits and is not represented in Figure 17.

[0066] One of ordinary skill in the art will recognize that one of the other computed values P'''(7) through P'''(1) may be selected, divided by 16, and used as the result of the interpolation. Using values corresponding to different vertices of the final sub-cube to generate the interpolation result will bias the result of the interpolation differently. This is a consideration in selecting which vertex value of the final sub-cube generated will be divided by 16 to generate the result of 40 the interpolation.

[0067] Shown in Figure 18 is a hardware implementation of a pruned tetrahedral interpolator 1000 for conversion of 45 input color space values 10 to a component of output color space values 11. This same hardware could be used repetitively for an additional (D-1) passes to generate the remaining D-1 components of the output color space value 11. Or, there could be an additional (D-1) replications of part of the hardware implementation shown in Figure 18 to generate each of D components simultaneously. The hardware used for generating multiplexer control inputs could be used for each of the D replications. The pruned tetrahedral interpolator 1000 shown in Figure 18 corresponds to d=3 and n=4 for the input color space value 10. The pruned tetrahedral interpolator 1000 implements the diagrammatic representation of pruned tetrahedral interpolation shown in Figure 15. The implementation of pruned tetrahedral interpolator 1000 requires considerably less hardware than the implementation of tetrahedral interpolator 900.

[0068] For d=3, bitwise AND blocks 1001a through 1001k each perform a bit by bit AND operation on three bit input 50 quantities to generate 3 bit output quantities. Likewise, for d=3, bitwise OR blocks 1002a through 1002k each perform a bit by bit OR operation on 3 bit input quantities to generate 3 bit output quantities. Each of the 3 bit outputs of bitwise AND blocks 1001a-1001g and bitwise OR blocks 1002a-1002g is used to control the selection of one of eight, 8 bit quantities in the corresponding of multiplexers 1003a through 1003n. The outputs of each of multiplexers 1003a-1003o are connected to the inputs of adders 1004a-1004h.

[0069] An interpolation operation is performed using pruned tetrahedral interpolator 1000 by supplying the vertex values selected using higher order bits 10a to the inputs of multiplexers 1003a-1003o and supplying the computed values of v[i] to the inputs of bitwise OR blocks 1002a-1002k and bitwise AND blocks 1001a-1001k as shown in Figure 18. In 55

addition, a vertex value selected using higher order bits 10a is supplied to the input of adder 1004h. Using the computed values of v[i], the bitwise AND blocks 1001a-1001k and the bitwise OR blocks 1002a-1002k compute the values input to the control inputs of multiplexers 1003a-1003n. Multiplexer 1003o uses v[3] directly. The values selected by multiplexers 1003a-1003o are those necessary to compute the interpolation result according to the diagrammatic representation of pruned tetrahedral interpolation shown in Figure 15. The vertex values selected by multiplexers 1003a-1003o are sent to the inputs of adders 1004a-1004h for summation. The output of the last adder in the chain of additions is divided by 16 and used as the result of the interpolation. The division by 16 is accomplished by bit shifting and is not shown in Figure 18.

[0070] The pruned tetrahedral interpolator shown in Figure 18 is implemented for d=3 and n=4. For some applications, less than four iterations of interpolation may be sufficient. Other applications may require more than four iterations of interpolation. The hardware implementation of pruned tetrahedral interpolation for d=3 and n=1 would use only a single adder and a single multiplexer to generate P'(0), as shown in Figure 18. The hardware implementations for n=2 and n=3 to generate, respectively, P"(0) and P'''(0) as shown in Figure 18, require more bitwise AND blocks, bitwise OR blocks, multiplexers, and adders.

[0071] Shown in Figure 19 is a high level flow diagram of a generalized method for performing pruned tetrahedral interpolation. For this method, the input color space values 10 are formed of d components. Each of the d components is partitioned into a set of higher order bits 10a and lower order bits 10b. Each of the d sets of lower order bits 10b is formed of n bits. The d sets of lower order bits are each designated as 1b₁, 1b₂, 1b₃, ..., 1b_d. The bit position of each of the lower order bits is designated from the most significant bit to the least significant bit by a value of i ranging, correspondingly from n-1 to 0. First, 2ⁿ⁻² values are computed 1100 according to the types of bitwise AND and bitwise OR operations shown in Figure 15 with $v[i]=2^{d-1} \times 1b_1[i] + 2^{d-2} \times 1b_2[i] + 2^{d-3} \times 1b_3[i] + \dots + 2^{d-d} \times 1b_d[i]$ for i ranging from n-1 to 0. Next, the minimum of 2ⁿ and 2^d of interpolation data values are selected 1101 using unique ones of the 2ⁿ⁻² values computed in step 1100, a value of v[n-1], and higher order bits 10a. Then, the sum of the interpolation data values selected in step 1101 is computed 1102.

[0072] The method shown in Figure 19 is for n iterations of pruned tetrahedral interpolation. The sum computed in step 1102 is divided by 2ⁿ to generate the result. This division is not shown in Figure 19.

[0073] Shown in Figure 20 is a high level flow diagram of a method implemented in software to perform pruned tetrahedral interpolation. First, a determination 1200 is made if any one of the components of the input color space value (a, b, c) 10 corresponds to a location on an outer boundary of the cubic lattice 1. This is the case if any one or more of the components of the input color space value has a value of FF hexadecimal. If this is the case, then, for purposes of generating the index into the cubic lattice 1 to retrieve the necessary vertex values, the components of the input color space value 10 which have a value of FF hexadecimal are assigned 1201 a value of 100 hexadecimal.

[0074] Next, the offsets from the origin of the cube in cubic lattice 1 accessed using higher order bits 10a and the subcubes generated during the pruned tetrahedral interpolation are computed 1202 using the relationships shown in the diagrammatic representation of the pruned tetrahedral interpolation of Figure 15. Then, the indices used to access the values corresponding to the required vertices of the selected cube in cubic lattice 1 in a look-up table are computed 1203. Finally, the values for each component of the output color space value (x, y, z) 11 are computed 1204. Provided in Figure 40 of this specification is the code of an implementation in C, for n=4, of the high level method of pruned tetrahedral interpolation shown in Figure 20.

[0075] It is also possible to implement pruned tetrahedral interpolation in hardware. The computational efficiencies which existed in the software implementation of pruned tetrahedral interpolation are present in the hardware implementation as reduced hardware requirements. As previously mentioned, shifts and concatenations are implemented without requiring additional hardware elements. The hardware functional blocks required to perform the pruned tetrahedral interpolation include adders, AND gates, OR gates, and multiplexers. With D dimensions in the output color space, d dimensions in the input color space, and n bits representing each group of lower order bits 10b of the input color space value, the requirements of the hardware implementation of the pruned subdivision interpolation can be computed as:

$$\# \text{ of Adders} = D \times 2^n \quad \text{eqn. 11}$$

$$\# \text{ of Multiplexers} = D \times (2^n - 1) \quad \text{eqn. 12}$$

$$\# \text{ of Additional AND/OR Gates} = D \times \left\{ 2 \times \left[\sum_{i=1}^{n-1} (2^i - 1) \right] \right\} \quad \text{eqn. 13}$$

[0076] To generate the gate level design necessary to implement the pruned radial interpolation in hardware, a com-

monly used hardware description language, such as VHDL, may be used. Included in Figure 41 is a listing of the VHDL code which can generate a hardware implementation of pruned radial interpolation.

[0077] As previously discussed, radial interpolation can result in print artifacts in conversions between the RGB and CMY color spaces for certain input color space values. Because of these print artifacts, a tetrahedral interpolation may 5 yield more desirable results. To reduce complexity, the tetrahedral interpolation can be implemented using pruned tetrahedral interpolation, although this interpolation technique is still more computationally intensive than the radial interpolation.

[0078] However, for conversions between other color spaces (such as CieLab, LUV, or YCbCr) the radial interpolation 10 may be preferable because it yields adequate results and is very computationally efficient. Additionally, it is possible that radial interpolation may actually produce more pleasing results than tetrahedral interpolation in some cases. If the interpolation methods are implemented in software, using alternative methods is easily done by calling different routines. However, implementing different interpolation methods in hardware can require separate logic. Because the separate hardware implementations of the two interpolation techniques are under utilized, this solution is expensive.

[0079] A common hardware implementation provides the capability for alternatively performing radial interpolation 15 and pruned tetrahedral interpolation with less hardware than a separate hardware implementation of these interpolation techniques. Shown in Figure 21 is a diagrammatic representation of a common radial interpolation and pruned tetrahedral interpolation implementation. As indicated in Figure 21, the interpolation technique performed is determined by the vertex values which are input to the hardware. Usually, the number of bits used to express each $v(i)$ term is fewer than the number of bits used to express each of the $P[v(i)]$ terms. Because of this, it is generally less complex to multiplex 20 the $v(i)$ terms prior to the memory access to retrieve the values associated with the $P[v(i)]$ terms.

[0080] It should be noted from Figure 21 that two of the vertex values used for both the radial interpolation and the pruned tetrahedral interpolation are the same for all values of n . Therefore, a common hardware implementation of radial interpolation and pruned tetrahedral interpolation requires an additional $2^n - 2$ multiplexers (each having d control bits) to be added to the hardware implementation of the pruned tetrahedral interpolation. Included in Figure 42 is the 25 VHDL code for a hardware implementation of common radial interpolation and pruned tetrahedral interpolation.

[0081] Shown in Figure 22 is a hardware implementation of a common pruned radial and pruned tetrahedral interpolator 1300. The hardware implementation of common pruned radial and pruned tetrahedral interpolation is similar to that of pruned tetrahedral interpolation. The difference is the addition of 14 multiplexers 1301a-1301n used to select the data to the control inputs of multiplexers 1302a-1302n. A single bit is used to control the selection of the data at the 30 inputs of multiplexers 1301a-1301n. The single bit controls whether the multiplexer control inputs to multiplexers 1302a-1302n are for pruned tetrahedral interpolation or for pruned radial interpolation. The multiplexer control input for multiplexers 1302a-1302n determines which of the interpolation data values are coupled to the adders. With the bit in the first of its two states, the hardware of Figure 22 performs as a pruned radial interpolator. With the bit in the second of its two states, the hardware of Figure 22 performs as a pruned tetrahedral interpolator.

[0082] The control bit for multiplexers 1301a-1301n is used to select between values of $v(i)$ and values computed 35 using bitwise OR blocks 1303a-1303k and bitwise AND blocks 1304a-1304k. Adders 1305a-1305o sum the outputs of multiplexers 1302a-1302o. By shifting bits, the resulting sum is divided by 16 (not shown in Figure 22) to generate the result.

[0083] The hardware of Figure 22 could be used repetitively for an additional ($D-1$) passes to generate the remaining 40 D-1 components of the output color space value 11. Or, there could be an additional ($D-1$) replications of part of the hardware implementation shown in Figure 22 to generate each of D components simultaneously. The hardware used for generating multiplexer control inputs could be used for each of the D replications. The common pruned radial and pruned tetrahedral interpolator 1300 shown in Figure 22 corresponds to $d=3$ and $n=4$ for the input color space value 10. The common pruned radial and pruned tetrahedral interpolator 1300 implements the diagrammatic representation of 45 common pruned radial and pruned tetrahedral interpolation shown in Figure 21.

[0084] In the interpolation process, the higher order bits 10a of an input color space value 10 form an index used to 50 access interpolation data values. The interpolation is performed using the lower order bits 10b of the input color space value. As previously mentioned, the accessed values correspond to the vertices of a cube in a cubic lattice 1. Depending upon the characteristics of the output color space, the values associated with the vertices of the accessed cube may vary at rates dependent upon the dimension of cubic lattice 1 or dependent upon the region of the cubic lattice 1 in which the selected cube is located. Because of this possibility, improved interpolation results may be produced by varying the interpolation resolution between values of the vertices throughout cubic lattice 1 corresponding to the varying 55 rates of change between vertex values. Adjusting the interpolation resolution based upon the location of the selected cube within cubic lattice 1 can be implemented by allowing the partitioning of the input color space value 10 into upper order bits 10a and lower order bits 10b to vary. In regions of cubic lattice 1 having high non-linear rates of change in the values of the vertices, the differences between the values of the vertices would be relatively large using a number of bits not adapted to the color space characteristics to represent each component of the upper order bits 10a. To reduce the values between the vertices, a larger number of bits are used to represent the components of the higher order bits

10a. Consequently, in these regions a smaller number of bits are used to represent the lower order bits 10b. For regions of cubic lattice 1 having lower rates of change or more linear rates of change in the values of the vertices, the differences between the values of the vertices using a number of bits not adapted to the color space characteristics to represent each component of the upper order bits 10a would be relatively small or relatively linear. To increase the values between the vertices a smaller number of bits are used to represent the components of the higher order bits 10a.

5 [0085] To implement an interpolation technique which can support a varying interpolation resolution over the output color space represented by cubic lattice 1, the interpolation technique must accommodate the changing number of bits used to represent the components of the lower order bits 10b. To accomplish this, a value, (n, p, q), is defined so that n bits are used to interpolate between lattice points in the "a" dimension, p bits are used to interpolate between lattice 10 points in the b dimension, and q bits are used to interpolate between lattice points in the c dimension. It is possible to constrain the values of each of the n, p, and q so that they are fixed over the entirety of cubic lattice 1 or to permit each of the values of the n, p, and q to vary independently or co-dependently throughout regions of the cubic lattice 1.

10 [0086] Shown in Figure 23a through Figure 23e is a graphical representation of a non-symmetric radial interpolation process that uses 4 bits to represent the a₁ component, 3 bits to represent the b₁ component, and 2 bits to represent the c₁ component of the lower order bits 10b of the input color space value. The first iteration of the cube subdivision is the selection of a sub-cube occupying one half of the cube selected by the higher order bits 10a using the bit from a₁ corresponding to the i=3 position. There are no bits present for the b₁ and the c₁ for the i=3 position. The second iteration of cube subdivision is the selection of a sub-cube occupying one fourth of the previous sub-cube using one bit each from a₁ and from b₁. The third iteration of cube subdivision is the selection of a sub-cube occupying one eighth of the previous sub-cube using one bit each from a₁, b₁, and c₁. The fourth iteration of cube subdivision is also the selection of a sub-cube occupying one eighth of the previous sub-cube using one bit each from a₁, b₁, and c₁. As can be seen from this, the number of bits of the components of the lower order bits 10b available to generate the sub-cube determines the fraction of the cube used to generate the sub-cube occupied by the sub-cube.

15 [0087] Shown in Table 4 and equations 14 through 19 are the relationships necessary to calculate the values of the sub-cube vertices for each iteration of radial sub-cube generation. Equations 14 through 19 generate the values used in the relationships shown in table 4 so that the correct sub-cube vertex values will be generated with or without the corresponding bits of a₁, b₁, and c₁, present for that iteration of sub-cube generation. If for a given iteration of sub-cube generation, a bit in any one or more of a₁, b₁, or c₁ is not present, equation 19 will generate the number of the vertex of the cube used in generating a vertex value of the sub-cube to compensate for the missing bit(s).

20 $30 \quad \text{Mask}_a = (2^n - 1) \quad \text{eqn. 14}$

$35 \quad \text{Mask}_b = (2^p - 1) \quad \text{eqn. 15}$

$40 \quad \text{Mask}_c = (2^q - 1) \quad \text{eqn. 16}$

$45 \quad m[i] = (4 \times \text{Mask}_a[i]) + (2 \times \text{Mask}_b[i]) + \text{Mask}_c[i] \quad \text{eqn. 17}$

$50 \quad v[i] = m[i] \& ((4 \times a[i]) + (2 \times b[i]) + c[i]) \quad \text{eqn. 18}$

$55 \quad f(N,i) = v[i] \& (N \& \sim m[i]) \quad \text{eqn. 19}$

Table 4

$P'[7] = \{P[7] + P[f(7,i)]\} \div 2$

$P'[6] = \{P[6] + P[f(6,i)]\} \div 2$

$P'[5] = \{P[5] + P[f(5,i)]\} \div 2$

$P'[4] = \{P[4] + P[f(4,i)]\} \div 2$

$P'[3] = \{P[3] + P[f(3,i)]\} \div 2$

$P'[2] = \{P[2] + P[f(2,i)]\} \div 2$

$P'[1] = \{P[1] + P[f(1,i)]\} \div 2$

$P'[0] = \{P[0] + P[f(0,i)]\} \div 2$

[0088] Shown in Figure 24 is a graphical representation of the generation of a sub-cube 1400 from a cube 1401 using non-symmetric radial sub-cube generation. On the particular iteration of the sub-cube generation shown in Figure 15, the bit corresponding to the iteration for components b_1 and c_1 is not present. Therefore, the relationships in table 4, with $f(N,i)$ calculated for the bits of b_1 and c_1 corresponding to the iteration not present, dictate that the vertex values 5 $P'[7]$, $P'[6]$, $P'[5]$, and $P'[4]$ are calculated as the average of the values of the two vertices vertically aligned with each of $P[7]$, $P[6]$, $P[5]$, and $P[4]$. The computation of vertex values for other combinations in which each of a_1 , b_1 , or c_1 are present or not present is handled analogously by equation 19.

[0089] Shown in Figure 25 is a diagrammatic representation of the non-symmetric pruned radial interpolation computation. The number of computations required to perform the non-symmetric pruned radial interpolation is computed as:

10

$$\# \text{ ALU operations} = [n \times [(4 \times d) - 1]] + [3 \times [2^{n-1} \times \sum_{i=1}^{n-1} (1 - 2^{-i})]] + [D \times (2^n + 1)] - 1 \quad \text{eqn. 20}$$

15

In Figure 25, 2^n memory references are shown. However, a selected cube has a maximum of 2^d vertex values, where d is the number of dimensions of the input color space value 10. Therefore, some of the 2^n (16 values for n=4) values shown at the inputs to the diagram of Figure 25 are redundant. It follows that for 2^n greater than 2^d , the number of memory accesses performed can be limited to the number of vertices in the cube for each dimension of the output of color 20 space. Therefore, the number of memory accesses required for D dimensions in the output color space is:

$$\# \text{ Memory accesses} = D \times \min(2^d, 2^n) \quad \text{eqn. 21}$$

[0090] Shown in Figure 26 is a high level flow diagram of a generalized method for performing non-symmetric pruned radial interpolation. For this method, the input color space values 10 are formed of d components. Each of the d components is partitioned into a set of higher order bits 10a and lower order bits 10b. The d sets of lower order bits are each designated as $1b_1, 1b_2, 1b_3, \dots, 1b_d$. Each of the d sets of lower order bits 10b is formed from respectively, of $n_1, n_2, n_3, \dots, n_d$ bits. The bit position of each of the d sets of lower order bits is designated from the most significant bit to the least significant bit by corresponding values of $i_1, i_2, i_3, \dots, i_d$ each ranging, correspondingly, from n_1-1 to 0, n_2-1 to 0, n_3-1 to 0, ..., n_d-1 to 0. First, a set of 2^{n-n-1} values is computed 1500 using $f(N, i) = v[i](N \& \sim m[i])$, where $m[i] = 2^{d-1} \times \text{Mask}_1[i] + 2^{d-2} \times \text{Mask}_2[i] + 2^{d-3} \times \text{Mask}_3[i] + \dots + 2^{d-d} \times \text{Mask}_d[i]$. For this computation the values of $\text{Mask}_k = 2^k - 1$ are each computed for a value k selected from $n_1, n_2, n_3, \dots, n_d$ with the value of j corresponding to the value of the subscript of the selected one of $n_1, n_2, n_3, \dots, n_d$. The values of j range from 1 to d. The values of $v[i]$ are computed as $m[i] \& (2^{d-1} \times 1b_1[i] + 2^{d-2} \times 1b_2[i] + 2^{d-3} \times 1b_3[i] + \dots + 2^{d-d} \times 1b_d[i])$ for values of i ranging from $n-1$ to 0, where n equals the greatest of $n_1, n_2, n_3, \dots, n_d$. The value of N corresponds to the vertex numbers ranging from 1 to 2^d . Next, a number of interpolation data values, equal to the minimum of 2^n and 2^d are selected 1501 using the unique values in the set of 2^{n-n-1} values, values of $v[i]$ for i ranging from $n-1$ to 0, and the d sets of higher order bits. Finally, the selected interpolation data values are added 1502 to generate a sum. To avoid roundoff error the generated sum is divided by 2^n . This 40 step is not shown in Figure 26.

[0091] Shown in Figure 27 is a high level flow diagram of method implemented in software to perform non-symmetric radial interpolation. First, the mask values are generated 1600 for each component of the input color space value 10. Next, a determination 1601 is made if any one of the components of the input color space value (a, b, c) 10 corresponds to a location on an outer boundary of the cubic lattice 1. This is the case if any one or more of the components of the 45 input color space value has a value of FF hexadecimal. If this is the case, then, for purposes of generating the indices to retrieve the necessary vertex values, the components of the input color space value 10 which have a value of FF hexadecimal are assigned 1602 a value of 100 hexadecimal. Then, the values of each of the $m[i]$ and $v[i]$ are computed 1603. Next, the indices used to access each of the vertex values used for the interpolation are computed 1604. Finally, each of the components of the output color space value are computed 1605 using the values accessed by the indices 50 computed in step 1604.

[0092] Included in Figure 43 is a listing in C of the code for a software implementation of non-symmetric radial interpolation. For the non-symmetric radial interpolation, the computed indices correspond to offsets from the origin of the cube selected by the higher order bits 10a. Because of the changing resolution used throughout the output color space, the values of the vertices for three cubes (one cube for each dimension of the output color space) selected by the higher 55 order bits 10a is passed into the routine of Figure 43 each time color space conversion is performed on input color space value 10. This is different than the code for the pruned radial and pruned tetrahedral interpolation in which the color table is passed as an array into the routine and indices into this table are computed in the routine. It is also possible to implement non-symmetric radial interpolation in hardware. As previously mentioned, shifts and concatenations

are implemented without requiring additional hardware elements. The hardware functional blocks required to perform the non-symmetric radial interpolation include adders, AND gates, OR gates, and multiplexers. With D dimensions in the output color space, d dimensions in the input color space, and n bits representing the maximum number of bits used to represent one of the components of the input color space value, the requirements of the hardware implementation of the non-symmetric radial interpolation can be computed as:

5

$$\# \text{ of Adders} = D \times [(2^n - 1) + 1] \quad \text{eqn. 22}$$

10

$$\# \text{ of Multiplexers} = D \times (2^n - 1) \quad \text{eqn. 23}$$

15

$$\# \text{ of additional AND/OR Gates} = (3 \times d) \times ((2^n - 1) \times \sum_{i=1}^{n-1} (1 - 2^{-i})) \quad \text{eqn. 24}$$

[0093] To generate the gate level design necessary to implement the non-symmetric pruned radial interpolation in hardware, a commonly used hardware description language, such as VHDL, may be used. Included in Figure 44 is a listing of the VHDL code which can generate a hardware implementation of non-symmetric pruned radial interpolation.

[0094] Shown in Figure 28 is a hardware implementation of a non-symmetric pruned radial interpolator 1700. Control input computation blocks 1701a-1701k compute the values used by the control inputs of multiplexers 1702a-1702o that are coupled to control input computation blocks 1701a-1701k. Each of control input computation blocks 1701a-1701k performs the computations of equations 14-19 on the input to that control input computation block. As shown in Figure 28, some of the control inputs of multiplexers 1702a-1702o use values of v[i]. Adders 1703a-1703o sum the outputs of multiplexers 1702a-1702o. This sum is divided by 2^n through bit shifting (not shown in Figure 28) to generate the interpolation result.

[0095] The hardware of Figure 28 could be used repetitively for an additional (D-1) passes to generate the remaining D-1 components of the output color space value 11. Or, there could be an additional (D-1) replications of part of the hardware implementation shown in Figure 28 to generate each of D components simultaneously. The hardware used for generating multiplexer control inputs could be used for each of the D replications. The non-symmetric pruned radial interpolator 1700 shown in Figure 28 corresponds to d=3 and the maximum one of $n_1, n_2, n_3, \dots, n_d$ equal to 4 for the input color space value 10. The non-symmetric pruned radial interpolator 1700 implements the diagrammatic representation of pruned tetrahedral interpolation shown in Figure 25.

[0096] Shown in Figure 29 is a high level flow diagram of a generalized method for performing non-symmetric radial interpolation using the equations of table 4. First, a first set of 2^d values are computed 1800 using $f(N, i) = v[i](N \& \neg m[i])$. Next, 2^d pairs of interpolation data values are selected 1801 with each pair formed from the interpolation data value selected using one of the first set of 2^d values and the interpolation data value corresponding the vertex designated by the value of N. Finally, 2^d sums are computed 1802 from the selected 2^d pairs of interpolation data values.

[0097] The method shown in Figure 29 is for a single iteration of non-symmetric radial interpolation. It should be recognized that further iterations would be performed by repeating the steps of Figure 29 with the successive sets of 2^d values computed using $f(N, i)$ for values of $v[i]$ and $m[i]$ corresponding to successively decremented values of i, selecting successive sets of 2^d pairs of values from the previously computed set of 2^d sums, and computing successive sets of 2^d sums from the successive sets of 2^d pairs of values. After performing n iterations, where n equals the greatest of $n_1, n_2, n_3, \dots, n_d$, one of the 2^d sums of the last set computed is divided by 2^n to generate the result. The division by 2 that could be performed after each iteration is delayed until after the final iteration to avoid round-off error.

[0098] Shown in Figure 30 is a hardware implementation of a non-symmetric radial interpolator 1900 for conversion of input color space values 10 to a component of output color space values 11. This same hardware could be used repetitively for an additional (D-1) passes to generate the remaining D-1 components of the output color space value 11. Or, there could be an additional (D-1) replications of part of the hardware implementation shown in Figure 30 to generate each of D components simultaneously. The hardware used for generating multiplexer control inputs could be used for each of the D replications. The non-symmetric radial interpolator 1900 shown in Figure 30 corresponds to d=3 and n=4, where n equals the greatest of $n_1, n_2, n_3, \dots, n_d$, for the input color space value 10. The hardware implementation shown in Figure 30 implements equations 14-19 and the equations of table 4 for the generation of the sub-cube vertex values. The non-symmetric radial interpolator 1900 of Figure 30 is formed from a first, second, third, and fourth stage 1901-1904. Each of the four stages 1901-1904 includes 2^3 adders, one of which is labeled as 1905. Each of the four stages 1901-1904 further includes 2^3 multiplexers, one of which is labeled as 1906. Finally, each of the four stages 1901-1904 includes 2^3 control input computation blocks for performing the computations of equations 14-19 with the indicated inputs. One of these control input computations blocks is labeled as 1907.

[0099] Each of stages 1901-1904 performs an iteration of interpolation. Some interpolation applications may require that only a single iteration of interpolation be performed. For a single iteration of interpolation n=1. This corresponds to a hardware implementation of non-symmetric radial interpolator 1900 using only first stage 1901. An additional stage would be added for each additional iteration of interpolation required for the particular application.

- 5 [0100] Each of the inputs of the multiplexers in the first stage 1901 are connected to the eight vertex values selected using higher order bits 10a. The outputs of each multiplexer in first stage 1901 are connected to the first input of the corresponding adder. The second input of the adder is connected to the value corresponding to the number of the vertex equal to one of the inputs of the corresponding control input computation blocks. The output of each of the adders of the first stage 1901 is a vertex value of the first sub-cube. As previously mentioned, the division by two for each iteration of sub-cube generation is deferred until the last sub-cube is generated. The vertex values of the last sub-cube generated are divided by 2^n , where n corresponds to the number of stages in the non-symmetric radial interpolator. The inputs of each multiplexer for the second, third, and fourth stages 1902-1904 are coupled to the outputs of the adders of the previous stage.

- 10 [0101] The control input of each multiplexer is connected to the output of the corresponding control input computation block. The multiplexers used in the non-symmetric radial interpolator 1900 have the capability to select one of eight, 15 eight bit values using a three bit control input. For this d=3, n=4 implementation of non-symmetric radial interpolator 1900, each of the inputs to the control input computation blocks is a 3 bit quantity. The output of each of the control input computation blocks is a 3 bit quantity. The adders associated with each multiplexer performs an addition of the selected eight bit values from each of the multiplexers.

- 20 [0102] Interpolation is performed by supplying the vertex values selected using higher order bits 10a to the multiplexer inputs of the first stage 1901. The inputs to the control input computation blocks are hardwired. The non-symmetric radial interpolator 1900 computes the vertex values for four iterations of sub-cube generation. The values P'''(7) through P'''(0) are the values of the vertices of the fourth sub-cube generated. In this embodiment, P'''(0) is selected, 25 divided by 16, and used as the result of the interpolation. The division by 16 is implemented by shifting bits and is not represented in Figure 30.

- [0103] One of ordinary skill in the art will recognize that one of the other computed values P'''(7) through P'''(1) may be selected, divided by 16, and used as the result of the interpolation. Using values corresponding to different vertices of the final sub-cube to generate the interpolation result will bias the result of the interpolation differently. This is a consideration in selecting which vertex value of the final sub-cube generated will be divided by 16 to generate the result of 30 the interpolation.

- [0104] Pruned tetrahedral interpolation can be adapted for implementation in a color space represented by a non-symmetric cubic lattice. As was the case for pruned tetrahedral interpolation, the vertices of the sub-cube generated are computed using the vertices of a tetrahedron partitioned from the cube used to generate the sub-cube. However, for some iterations of sub-cube generation, the corresponding bits of each of the a_1 , b_1 , and c_1 components may not be 35 present. For these cases, the computation of the sub-cube vertices must be modified to compensate for the missing corresponding bits in one or more of the a_1 , b_1 , and c_1 components.

- [0105] Shown in Table 5 and equations 25 and 26 are the relationships necessary to calculate the values of the sub-cube vertices for non-symmetric pruned tetrahedral sub-cube generation. Equations 25 and 26 along with equations 14 40 through 18 are used to generate the proper values so that in the relationships listed in table 5, the correct sub-cube vertex values will be generated with or without the bits of a_1 , b_1 , and c_1 corresponding to that iteration of sub-cube generation present. If for a given iteration of sub-cube generation, a bit in any one or more of a_1 , b_1 , and c_1 is not present, equation 25 and equation 26 will generate the number of the vertex of the cube used in generating a vertex value of the sub-cube so that compensation is made for the missing bit(s).

45 $g(N,i) = (v[i]) \sim m[i] \& N$ eqn. 25

$h(N,i) = (v[i])N$ eqn. 26

50

Table 5

$$P[7] = \{P[g(7,i)] + P[h(7,i)]\} \div 2$$

$$P[6] = \{P[g(6,i)] + P[h(6,i)]\} \div 2$$

$$P[5] = \{P[g(5,i)] + P[h(5,i)]\} \div 2$$

$$P[4] = \{P[g(4,i)] + P[h(4,i)]\} \div 2$$

EP 0 923 049 A2

Table 5 (continued)

$$P'[3] = \{P[g(3,i)] + P[h(3,i)]\} \div 2$$

$$P'[2] = \{P[g(2,i)] + P[h(2,i)]\} \div 2$$

$$P'[1] = \{P[g(1,i)] + P[h(1,i)]\} \div 2$$

$$P'[0] = \{P[g(0,i)] + P[h(0,i)]\} \div 2$$

5

[0106] Although Figure 24 provides a graphical representation of the generation of a sub-cube 1400 from a cube 1401 using non-symmetric radial sub-cube generation, Figure 24 can also provide a graphical representation of non-symmetric pruned tetrahedral sub-cube generation. On the particular iteration of the sub-cube generation shown in Figure 24, the bit corresponding to the iteration for components b_1 and c_1 , is not present. Therefore, the relationships in table 5, with $g(N,i)$ and $h(N,i)$ calculated for the corresponding bits of b_1 and c_1 not present, dictate that the vertex values $P'[7]$, $P'[6]$, $P'[5]$, and $P'[4]$ are calculated as the average of the values of the two vertices vertically aligned with each of $P'[7]$, $P'[6]$, $P'[5]$, and $P'[4]$. The computation of vertex values for other combinations in which each of a_1 , b_1 , or c_1 are present or not present is handled analogously by equations 25 and 26.

[0107] Shown in Figure 31 is a diagrammatic representation of the non-symmetric pruned tetrahedral interpolation computation. The number of computations required to perform the non-symmetric pruned tetrahedral interpolation is computed as:

20

$$\# \text{ of ALU operations} = n \times [(4 \times d) - 1] + 4 \times \left\{ 2^{n-1} \times \sum_{i=1}^{n-1} (1 - 2^{-i}) \right\} \quad \text{eqn. 27}$$

25

$$\# \text{ of Memory accesses} = D(\min(2^d, 2^n)) \quad \text{eqn. 28}$$

[0108] Shown in Figure 32 is a high level flow diagram of a generalized method for performing non-symmetric pruned tetrahedral interpolation. For this method, the input color space values 10 are formed of d components. Each of the d components is partitioned into a set of higher order bits 10a and lower order bits 10b. The d sets of lower order bits are each designated as $1b_1, 1b_2, 1b_3, \dots, 1b_d$. Each of the d sets of lower order bits 10b is formed from, respectively, of $n_1, n_2, n_3, \dots, n_d$ bits. The bit position of each of the d sets of lower order bits is designated from the most significant bit to the least significant bit by corresponding values of $i_1, i_2, i_3, \dots, i_d$ each ranging, correspondingly, from n_1-1 to 0, n_2-1 to 0, n_3-1 to 0, ..., n_d-1 to 0.

[0109] First, a set of 2^{n-2} values is computed 2000 using $g(N, i) = (v[i] \sim m[i] \& N)$ and $h(N, i) = (v[i] | N)$, where $m[i] = 2^{d-1} \times \text{Mask}_1[i] + 2^{d-2} \times \text{Mask}_2[i] + 2^{d-3} \times \text{Mask}_3[i] + \dots + 2^{d-d} \times \text{Mask}_d[i]$. For this computation the values of $\text{Mask}_j = 2^{k-1}$ are each computed for a value k selected from $n_1, n_2, n_3, \dots, n_d$ with the value of j corresponding to the value of the subscript of the selected one of $n_1, n_2, n_3, \dots, n_d$. The values of j range from 1 to d . The values of $v[i]$ are computed as $m[i] \& (2^{d-1} \times 1b_1[i] + 2^{d-2} \times 1b_2[i] + 2^{d-3} \times 1b_3[i] + \dots + 2^{d-d} \times 1b_d[i])$ for values of i ranging from $n-1$ to 0, where n equals the greatest of $n_1, n_2, n_3, \dots, n_d$. The value of N corresponds to the vertex numbers ranging from 1 to 2^d . Next, a number of interpolation data values, equal to the minimum of 2^n and 2^d are selected 2001 using the unique values in the set of 2^{n-2} values, a value of $v[i]$ for i in the range from $n-1$ to 0, and the d sets of higher order bits. Finally, the selected interpolation data values are added 2002 to generate a sum. To avoid roundoff error the generated sum is divided by 2^n . This step is not shown in Figure 32.

[0110] Shown in Figure 33 is a high level flow diagram of a method implemented in software to perform non-symmetric pruned tetrahedral interpolation. First, the mask values are generated 2100 for each component of the input color space value 10. Next, a determination 2101 is made if any one of the components of the input color space value (a, b, c) 10 corresponds to a location on an outer boundary of the cubic lattice 1. This is the case if any one or more of the components of the input color space value has a value of FF hexadecimal. If this is the case, then, for purposes of generating the index into the cubic lattice 1 to retrieve the necessary vertex values, the components of the input color space value 10 which have a value of FF hexadecimal are assigned 2102 a value of 100 hexadecimal. Then, the values of each of the $m[i]$ and $v[i]$ are computed 2103. Next, the indices used to access each of the vertex values used for the interpolation are computed 2104 using $g(N, i)$ and $h(N, i)$. Finally, each of the components of the output color space value are computed 2105 using the values accessed by the indices computed in step 2104.

[0111] Included in Figure 45 is a listing of the code for a software implementation of non-symmetric pruned tetrahedral interpolation in C. For the non-symmetric pruned tetrahedral interpolation, the computed indices correspond to offsets from the origin of the cube selected by the higher order bits 10a. Because of the changing resolution used throughout the output color space, the values of the vertices for three cubes (one cube for each dimension of the output color

space) selected by the higher order bits 10a is passed into the routine of Figure 45 each time color space conversion is performed on an input color space value 10. This is different than the code for the pruned radial and pruned tetrahedral interpolation in which the color table is passed as an array into the routine and indices into this table are computed in the routine.

- 5 [0112] It is also possible to implement non-symmetric pruned tetrahedral interpolation in hardware. As previously mentioned, shifts and concatenations are implemented without requiring additional hardware elements. The hardware functional blocks required to perform the non-symmetric pruned tetrahedral interpolation include adders, AND gates, OR gates, and multiplexers. With D dimensions in the output color space, d dimensions in the input color space, and n bits representing the maximum number of bits used to represent one of the components of the input color space value,
10 the requirements of the hardware implementation of the non-symmetric pruned tetrahedral interpolation can be computed as:

$$\# \text{ of Adders} = D \times 2^n \quad \text{eqn. 29}$$

15 # of Multiplexers = $D \times (2^n - 1)$ eqn. 30

20 # of Additional AND/OR Gates = $(4 \times d) \times \left\{ 2^{n-1} \times \sum_{i=1}^{n-1} (1 \cdot 2^{-i}) \right\}$ eqn. 31

- [0113] To generate the gate level design necessary to implement the non-symmetric pruned tetrahedral interpolation in hardware, a commonly used hardware description language, such as VHDL, may be used. Included in Figure 46 is
25 a listing of the VHDL code which can generate a hardware implementation of non-symmetric pruned tetrahedral interpolation.

- [0114] Shown in Figure 34 is a hardware implementation of a non-symmetric pruned tetrahedral interpolator 2200. Control input computation blocks 2201a-2201v compute the values used by the control inputs of multiplexers 2202a-2202n. The control input computation blocks 2201a-2201v apply the functions of equations 25 and 26, as indicated in
30 Figure 34 to compute the control inputs for multiplexers 2202a-2202n. As shown in Figure 34, multiplexer 22020 uses a value of $v[i]$ for its control input. Each of multiplexers 2202a-22020 select from eight interpolation data values selected using higher order bits 10a. Adders 2203a-2203o sum the outputs of multiplexers 2202a-2202o. This sum is divided by 2^n through bit shifting (not shown in Figure 34) to generate the interpolation result.

- [0115] The hardware of Figure 28 could be used repetitively for an additional (D-1) passes to generate the remaining
35 D-1 components of the output color space value 11. Or, there could be an additional (D-1) replications of part of the hardware implementation shown in Figure 28 to generate each of D components simultaneously. The hardware used for generating multiplexer control inputs could be used for each of the D replications. The non-symmetric pruned tetrahedral interpolator 2200 shown in Figure 34 corresponds to $d=3$ and the maximum one of $n_1, n_2, n_3, \dots, n_d$ equal to 4 for the input color space value 10. The non-symmetric pruned tetrahedral interpolator 2200 implements the diagrammatic
40 representation of non-symmetric pruned tetrahedral interpolation shown in Figure 31. Shown in Figure 35 is a high level flow diagram of a generalized method for performing non-symmetric tetrahedral interpolation. First, a first and a second set of 2^d values are computed 2300 using, respectively $g(N,i)$ and $h(N,i)$. Next, 2^d pairs of interpolation data values are selected 2301 using the first and second set of values. Finally a set of 2^d sums are computed 2302 using the 2^d pairs of interpolation data values.

- 45 [0116] The method shown in Figure 35 is for a single iteration of non-symmetric tetrahedral interpolation. It should be recognized that further iterations would be performed by repeating the steps of Figure 35 with the successive sets of 2^d values computed using $g(N,i)$ and $h(N,i)$ for values of $v[i]$ and $m[i]$ corresponding to successively decremented values of i, selecting successive sets of 2^d pairs of values from the previously computed set of 2^d sums, and computing successive sets of 2^d sums from the successive sets of 2^d pairs of values. After performing n iterations, where n equals the
50 greatest of $n_1, n_2, n_3, \dots, n_d$, one of the 2^d sums of the last set computed is divided by 2^n (not shown in Figure 35) to generate the result. The division by 2 that could be performed after each iteration is delayed until after the final iteration to avoid round-off error.

- [0117] Shown in Figure 36 is a hardware implementation of a non-symmetric tetrahedral interpolator 2400 for conversion of input color space values 10 to a component of output color space values 11. This same hardware could be used
55 repetitively for an additional (D-1) passes to generate the remaining D-1 components of the output color space value 11. Or, there could be an additional (D-1) replications of part of the hardware implementation shown in Figure 36 to generate each of D components simultaneously. The hardware used for generating multiplexer control inputs could be used for each of the D replications. The non-symmetric tetrahedral interpolator 2400 shown in Figure 36 corresponds to $d=3$

and $n=4$, where n equals the greatest of $n_1, n_2, n_3, \dots, n_d$, for the input color space value 10. The hardware implementation shown in Figure 36 implements equations 25 and 26 and the equations of table 5 for the generation of the sub-cube vertex values. The non-symmetric tetrahedral interpolator 2400 of Figure 36 is formed from a first, second, third, and fourth stage 2401-2404. Each of the four stages 2401-2404 includes 2^3 adders, one of which is labeled as 2405. Each of the four stages 1901-1904 further includes 2×2^3 multiplexers, one of which is labeled as 2406. The multiplexers for each stage are arranged into 2^3 pairs. Finally, each of the four stages 2401-2404 includes 2^3 control input computation blocks for computing $g(N,i)$ and 2^3 control input computation blocks for computing $h(N,i)$. One of the control input computation blocks for computing $g(N,i)$ is labeled as 2407 and one of the control input computation blocks for computing $h(N,i)$ is labeled as 2408.

5 [0118] Each of stages 2401-2404 performs an iteration of interpolation. Some interpolation applications may require that only a single iteration of interpolation be performed. For a single iteration of interpolation $n=1$. This corresponds to a hardware implementation of non-symmetric tetrahedral interpolator 2400 using only first stage 2401. An additional stage would be added for each additional iteration of interpolation required for the particular application.

10 [0119] Each of the inputs of the multiplexers in the first stage 2401 are connected to the eight vertex values selected using higher order bits 10a. The outputs of each pair of multiplexers in stages 2401-2404 are connected to, respectively, the first and second inputs of the corresponding adder. The output of each of the adders of the first stage 1901 is a vertex value of the first sub-cube. As previously mentioned, the division by 2 for each iteration of sub-cube generation is deferred until the last sub-cube is generated. The vertex values of the last sub-cube generated are divided by 2^n (not shown in Figure 36), where n corresponds to the number of stages in the non-symmetric tetrahedral interpolator. The 15 inputs of each multiplexer for the second, third, and fourth stages 2402-2404 are coupled to the outputs of the adders of the previous stage.

20 [0120] The control input of each multiplexer is connected to the output of the corresponding control input computation block. The multiplexers used in the non-symmetric tetrahedral interpolator 2400 have the capability to select one of eight, eight bit values using a three bit control input. For this $d=3$ implementation of non-symmetric tetrahedral interpolator 2400, each of the inputs to the control input computation blocks is a 3 bit quantity. The output of each of the control input computation blocks is a 3 bit quantity. The adders associated with each multiplexer performs an addition of the selected eight bit values from each of the multiplexers.

25 [0121] Interpolation is performed by supplying the vertex values selected using higher order bits 10a to the multiplexer inputs of the first stage 2401. The inputs to the control input computation blocks are hardwired. The non-symmetric tetrahedral interpolator 2400 computes the vertex values for four iterations of sub-cube generation. The values $P'''(7)$ through $P'''(0)$ are the values of the vertices of the fourth sub-cube generated. In this embodiment, $P'''(0)$ is selected, divided by 16, and used as the result of the interpolation. The division by 16 is implemented by shifting bits and is not represented in Figure 36.

30 [0122] One of ordinary skill in the art will recognize that one of the other computed values $P'''(7)$ through $P'''(1)$ maybe selected, divided by 16, and used as the result of the interpolation. Using values corresponding to different vertices of the final sub-cube to generate the interpolation result will bias the result of the interpolation differently. This is a consideration in selecting which vertex value of the final sub-cube generated will be divided by 16 to generate the result of the interpolation.

35 [0123] A common hardware implementation of non-symmetric radial interpolation and non-symmetric pruned tetrahedral interpolation is possible. As can be seen from the diagrammatic representations of the non-symmetric radial interpolation and the non-symmetric pruned tetrahedral interpolation in Figure 25 and Figure 31, respectively, a common hardware implementation could be accomplished by multiplexing the indices used to access the input vertex values. Included in Figure 47 is a listing of the VHDL code which can generate a common hardware implementation of 40 non-symmetric radial and non-symmetric pruned tetrahedral interpolation.

45 [0124] Shown in Figure 37 is a hardware implementation of a common non-symmetric pruned radial and non-symmetric pruned tetrahedral interpolator 2500. The hardware implementation of common non-symmetric pruned radial and non-symmetric pruned tetrahedral interpolation incorporates the control input computation blocks of the non-symmetric pruned radial 1700 and non-symmetric pruned tetrahedral 2200 interpolators. Multiplexers 2501a-2501n are used to select the data to the control inputs of multiplexers 2502a-2502n. A single bit is used to control the selection of 50 the data at the inputs of multiplexers 2501a-2501n. The single bit controls whether the multiplexer control inputs to multiplexers 2502a-2502n are for non-symmetric pruned tetrahedral interpolation or for non-symmetric pruned radial interpolation. The multiplexer control input for multiplexers 2502a-2502n determines which of the interpolation data values are coupled to the adders. With the bit in the first of its two states, the hardware of Figure 37 performs as a non-symmetric pruned radial interpolator. With the bit in the second of its two states, the hardware of Figure 37 performs as a non-symmetric pruned tetrahedral interpolator. Adders 2503a-2503o sum the outputs of multiplexers 2502a-2502o. By shifting bits, the resulting sum is divided by 16 (not shown in Figure 22) to generate the result. Control input computation blocks 2504a-2504k implement equations 14-19 and control input computation blocks 2505a-2505v implement equations 25 and 26.

[0125] The hardware of Figure 37 could be used repetitively for an additional (D-1) passes to generate the remaining D-1 components of the output color space value 11. Or, there could be an additional (D-1) replications of part of the hardware implementation shown in Figure 37 to generate each of D components simultaneously. The hardware used for generating multiplexer control inputs could be used for each of the D replications. The common non-symmetric pruned radial and non-symmetric pruned tetrahedral interpolator 2500 shown in Figure 37 corresponds to d=3 and n=4 for the input color space value 10.

[0126] It should be recognized that for each of the disclosed hardware embodiments of interpolators, computations are required to supply the multiplexer control inputs. These computations may be implemented in dedicated hardware or performed using a microprocessor under software control. Using a microprocessor to compute the multiplexer control 10 inputs results in a hardware savings at the expense of increasing the time required to perform the multiplexer control input computations.

[0127] Although several embodiments of the inventions have been illustrated, and their forms described, it is readily apparent to those of ordinary skill in the art that various modifications may be made therein without departing from the spirit of the invention or from the scope of the appended claims.

15

Claims

1. A pruned tetrahedral interpolator (1000) for interpolating between interpolation data values using input data values (10) each having d components to generate output data values (11), said d components represented by d sets of bits partitioned to form d sets of lower order bits (10b) with each of said d sets of lower order bits (10b) having n of said bits, said pruned tetrahedral interpolator (1000) comprising:

20 a first set of 2^{n-1} multiplexers (1003a-1003o) each configured for receiving one of a set of control inputs and having a multiplexer output, each of said multiplexers (1003a-1003o) of said first set for selecting from said interpolation data values responsive to said one of said set of control inputs; and
25 a means for adding (1004a-1004o) configured for receiving said multiplexer output of said set of multiplexers (1003a-1003o).

30

2. The pruned tetrahedral interpolator (1000) as recited in claim 1, wherein:

35

each of said multiplexers (1003a-1003o) of said first set includes 2^d of multiplexer inputs for receiving 2^d of said interpolation data values;

said means for adding (1004a-1004o) includes 2^{n-1} adders (1004a-1004o) each having a first input, a second input, and an output;

35 2^{n-1} of said adders (1004a-1004o) each have said first input and said second input coupled to one of said multiplexer output of said multiplexers (1003a-1003o) of said first set;

said first input of one of said adders (1004a-1004o) includes a configuration for receiving one of said interpolation data values selected using d sets of higher order bits (10a) partitioned from said d sets of bits and said second input couples to one of said multiplexer output of said multiplexers (1003a-1003o) of said first set;

40 $2^n \cdot (2^{n-1}-1)$ of said adders (1004a-1004o) each have said first input and said second input coupled to said output of another of said adders (1004a-1004o); and

said first set includes a partitioning into n groups of said multiplexers (1003a-1003o) each designated by a value of i ranging from n-1 to 0, each of said n groups of said first set having, respectively, 2^i of said multiplexers (1003a-1003o) from said first set.

45

3. The pruned tetrahedral interpolator (1000) as recited in claim 2, further comprising:

50

a second set of 2^{n-1} -1 bitwise AND blocks (1001a-1001k) each having a first input, a second input, and an output, said second set partitioned into n-1 groups of said bitwise AND blocks (1001a-1001k) each designated by said value of i ranging from n-1 to 1, each of said n-1 groups of said second set having, respectively, 2^{i-1} of said bitwise AND blocks (1001a-1001k), each of said output of said bitwise AND blocks (1001a-1001k) from said n-1 groups of said second set coupled to one of said multiplexers (1003a-1003o) in one of said n-1 groups of said first set having the corresponding of said value of i to supply one of said set of control inputs; and

a third set of 2^{n-1} -1 bitwise OR blocks (1002a-1002k) each having a first input, a second input, and an output, said third set partitioned into n-1 groups of said bitwise OR blocks (1002a-1002k) each designated by said value of i ranging from n-1 to 1, each of said n-1 groups of said third set having, respectively, 2^{i-1} of said bitwise OR blocks (1002a-1002k), each of said output of said bitwise OR blocks (1002a-1002k) from said n-1 groups of said third set coupled to one of said multiplexers (1003a-1003o) in one of said n-1 groups of said first set

having the corresponding of said value of i to supply one of said set of control inputs.

4. The pruned tetrahedral interpolator (1000) as recited in claim 3, wherein:

5 1b₁, 1b₂, 1b₃, ..., 1b_d designate said d sets of lower order bits (10b) with a bit position of each of said bits of said
d sets of lower order bits (10b) designated from a most significant of said lower order bits (10b) to a least sig-
nificant of said lower order bits (10b) by a value of k ranging, correspondingly, from n-1 to 0, computation of a
fourth set of n values, each designated as v[k], uses said d sets of lower order bits (10b) according to
v[k]=2^{d-1} × 1b₁[k] + 2^{d-2} × 1b₂[k] + 2^{d-3} × 1b₃[k] + ... + 2^{d-d} × 1b_d[k] for each of said values of k; and
10 said first input of each of said bitwise AND blocks (1001a-1001k) and said bitwise OR blocks (1002a-1002k) in
said n-1 groups corresponding to said value of i ranging from n-1 to 1 includes a configuration to receive one
of said fourth set of n values for k equal to n-1, said second input of each of said bitwise AND blocks (1001a-
1001k) and said bitwise OR blocks (1002a-1002k) in said n-1 groups corresponding to said value of i ranging
from n-1 to 2 each arranged for receiving one of a fifth set of values computed using bitwise AND operations
15 and bitwise OR operations upon ones of said fourth set of n values, said second input of each of said bitwise
AND blocks (1001a-1001k) and said bitwise OR blocks (1002a-1002k) in said n-1 groups corresponding to
said value of i equal to 1 includes a configuration to receive one of said fourth set of n values for k equal to n-2.

5. The pruned tetrahedral interpolator (1000) as recited in claim 4, wherein:

20 each of said interpolation data values includes D interpolation data value components;
each of said D interpolation data value components corresponds to a dimension of an output color space and
each of said d components of said input data values (10) corresponds to a dimension of an input color space;
a printing device includes said pruned tetrahedral interpolator (1000) for performing a color space conversion.
25

6. The pruned tetrahedral interpolator (1000) as recited in claim 5, wherein:

30 D equals 3, d equals 3, and n equals 4, said output color space includes a color space selected from the group
consisting of a RGB, a Lab, a XYZ, a HSV, a Luv, a HLS, and a CMY color space and said input color space
includes a color space selected from the group consisting of a RGB, a Lab, a XYZ, a HSV, a Luv, a HLS, and a
CMY color space.

7. The pruned tetrahedral interpolator (1000) as recited in claim 5, wherein:

35 D equals 4, d equals 3, and n equals 4, said output color space includes a CMYK color space and said input
color space includes a color space selected from the group consisting of a RGB, a Lab, a XYZ, a HSV, a Luv,
a HLS, and a CMY color space.

40 8. A tetrahedral interpolator (900) for interpolating between interpolation data values using input data values (10)
each having d components to generate output data values (11), said d components represented, correspondingly,
by d sets of bits each partitioned to form d sets of higher order bits (10a) and d sets of lower order bits (10b), said
d sets of higher order bits (10a) used for selecting 2^d of said interpolation data values, said tetrahedral interpolator
(900) comprising:

45 a set of 2 × 2^d multiplexers (906) with each of said multiplexers (906) having a multiplexer output and config-
ured to receive said 2^d of said interpolation data values, each of said multiplexers (906) for selecting a one of
said 2^d of said interpolation data values responsive to one of 2 × 2^d values determined from said d sets of lower
order bits (10b); and
50 a set of 2^d adders (905) with each of said adders (905) having a first input, a second input, and a output, each
of a first group of 2^d of said multiplexers (906) having the corresponding of said multiplexer output coupled to
one of said first input, each of a second group of 2^d of said multiplexers (906) having the corresponding of said
multiplexer output coupled to one of said second input, said set of 2 × 2^d multiplexers (906) and said set of 2^d
adders (905) forming a stage.

55 9. With interpolation data values for selection using input data values (10) each having d components, said d com-
ponents represented by d sets of bits each partitioned to form d sets of higher order bits (10a) and d sets of lower
order bits (10b) with each of said d sets of lower order bits (10b) having n bits, said d sets of lower order bits (10b)
designated as 1b₁, 1b₂, ..., 1b_d with the bit position of each bit of said d sets of lower order bits (10b) designated from

EP 0 923 049 A2

the most significant of said lower order bits (10b) to the least significant of said lower order bits (10b) by a value of i ranging, correspondingly, from n-1 to 0, a method of tetrahedral interpolation, comprising the steps of:

5 computing a first value according to $v[i] = 2^{d-1} \times 1b_1[i] + 2^{d-2} \times 1b_2[i] + \dots + 2^{d-d} \times 1b_d[i]$ for said value of i equal to (n-1);

computing a first set of AND values according to $v[n-1] \& k$, for said value of k ranging from 2^d-1 to 0;

computing a first set of OR values according to $v[n-1] | k$, for said value of k ranging from 2^d-1 to 0;

selecting a first set of 2^d pairs of said interpolation data values using said first set of said AND values and said first set of said OR values, each of said first set of 2^d pairs selected using one of said first set of said AND values and one of said first set of said OR values each computed using the same of said value of k; and
10 computing a first set of 2^d sums by summing each of said first set of 2^d pairs of said interpolation data values.

10. With interpolation data values for selection using input data values (10) each having d components, said d components represented by d sets of bits each partitioned to form d sets of higher order bits (10a) and d sets of lower order bits (10b) with each of said d sets of lower order bits (10b) having n bits, said d sets of lower order bits (10b) designated as $1b_1, 1b_2, \dots, 1b_d$ with the bit position of each bit of said d sets of lower order bits (10b) designated from the most significant of said lower order bits (10b) to the least significant of said lower order bits (10b) by a value of i ranging, correspondingly, from n-1 to 0, a method of pruned tetrahedral interpolation, comprising the steps of:

20 computing a first set of 2^n-2 values using bitwise AND operations and bitwise OR operations operating upon $v[i]$ equal to $2^{d-1} \times 1b_1[i] + 2^{d-2} \times 1b_2[i] + \dots + 2^{d-d} \times 1b_d[i]$ for said value of i ranging from (n-1) to 0;

selecting at least the minimum of 2^n and 2^d of said interpolation data values using ones of said first set of 2^n-2 values, $v[i]$ for one of said values of i equal n-1, and said d sets of higher order bits (10a); and

25 adding a second set of said interpolation data values formed from said interpolation data values from said step of selecting to generate a sum.

30

35

40

45

50

55

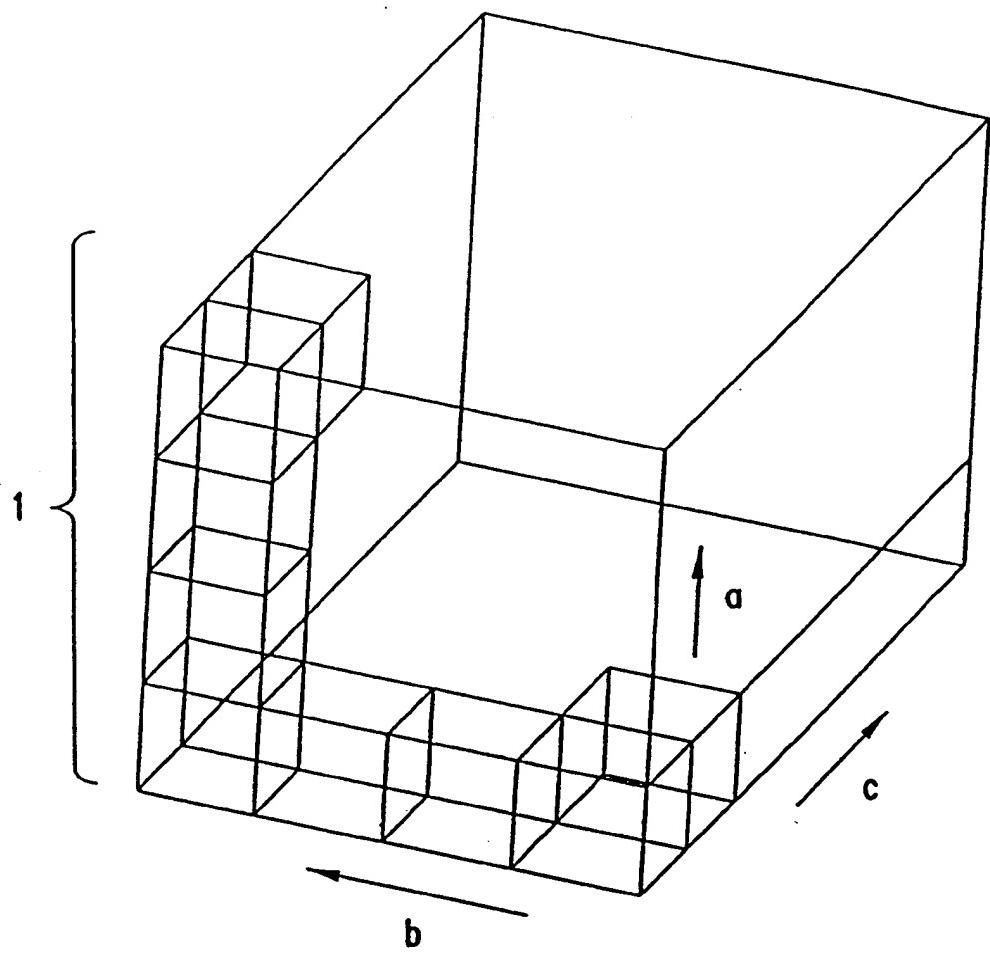


FIG.1A

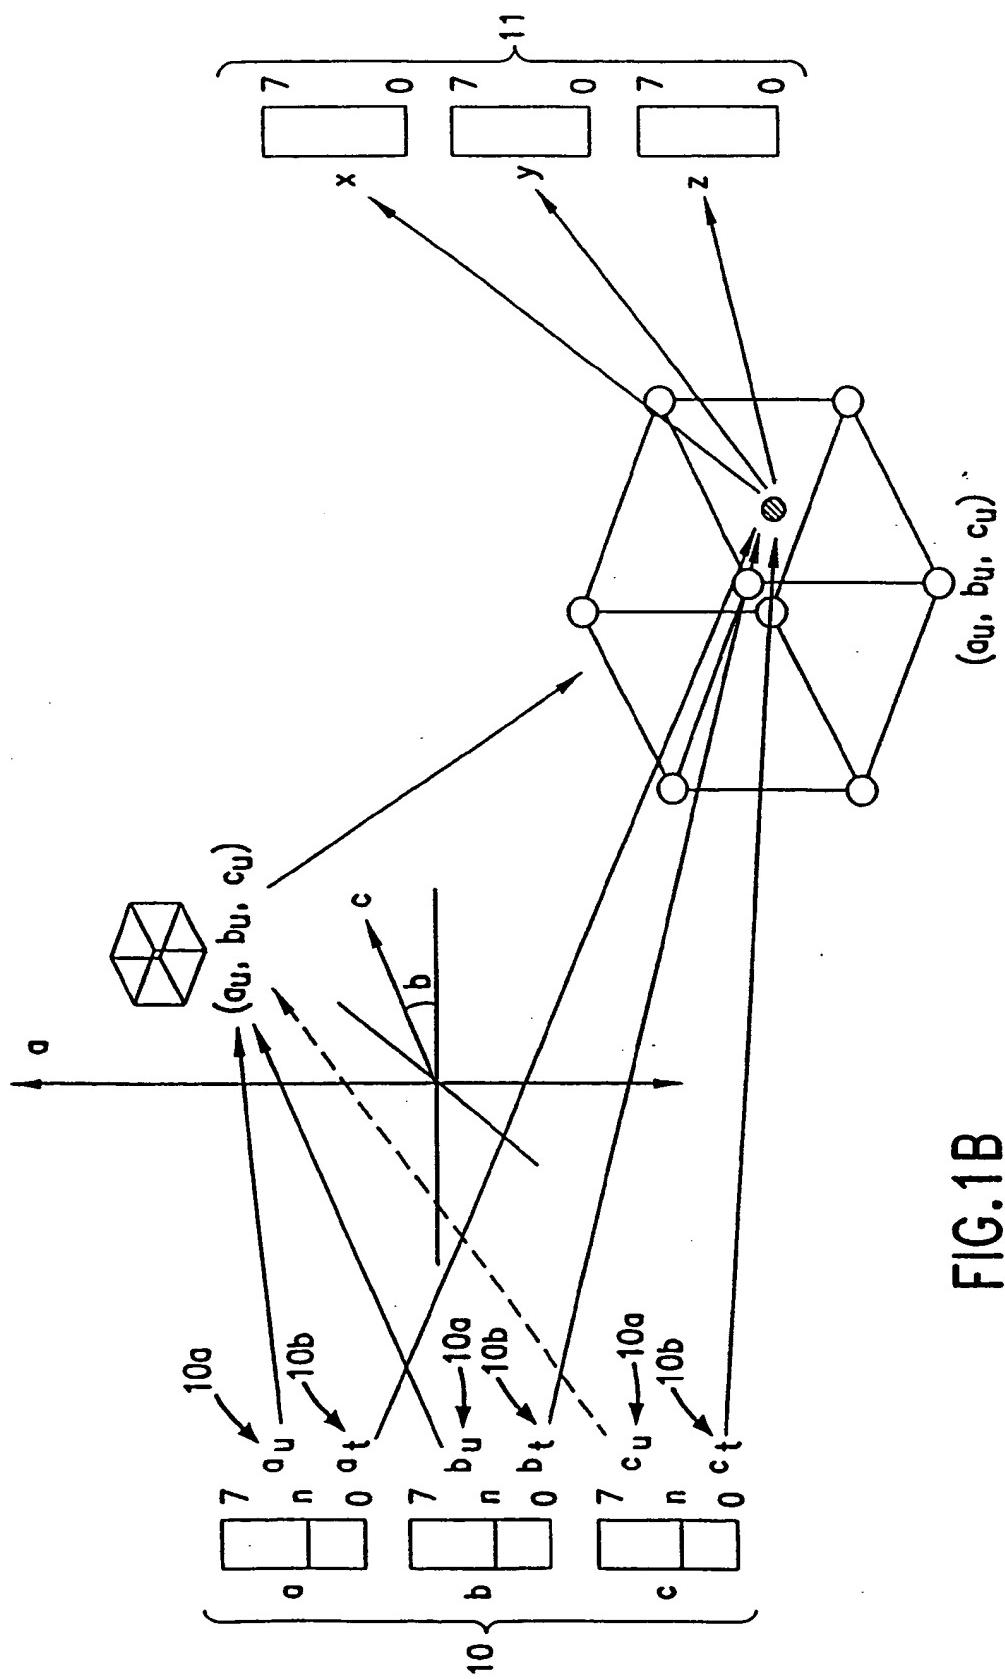


FIG. 1B

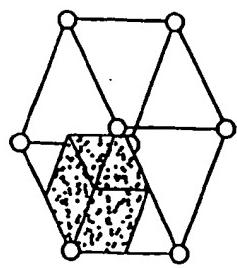


FIG.2D

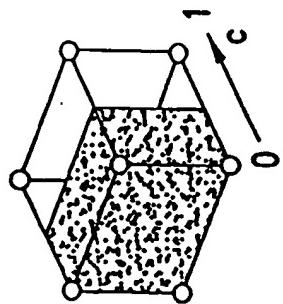


FIG.2C

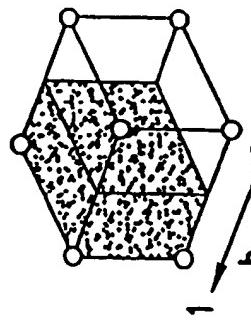


FIG.2B

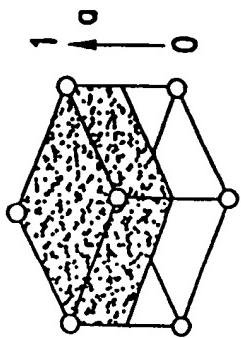


FIG.2A

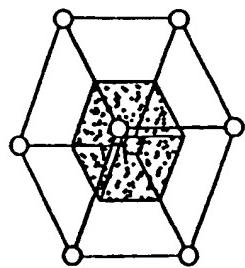


FIG.3D

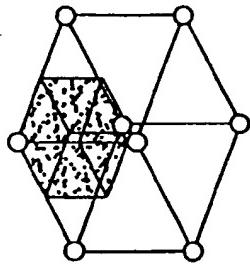


FIG.3H

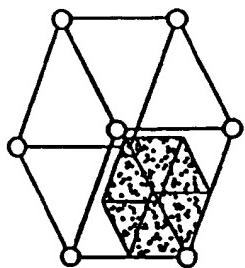


FIG.3C

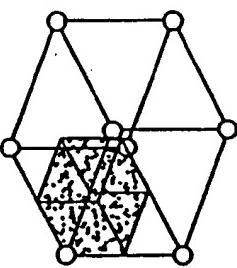


FIG.3G

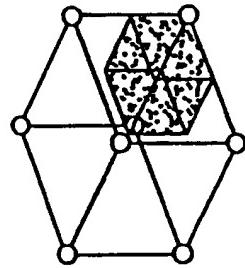


FIG.3B

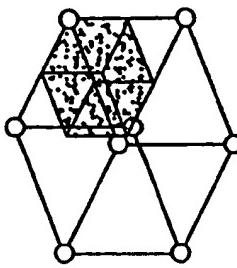


FIG.3F

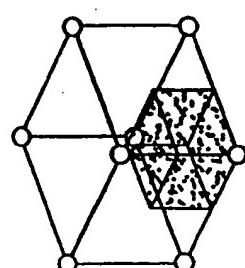


FIG.3A

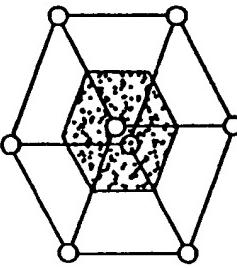


FIG.3E

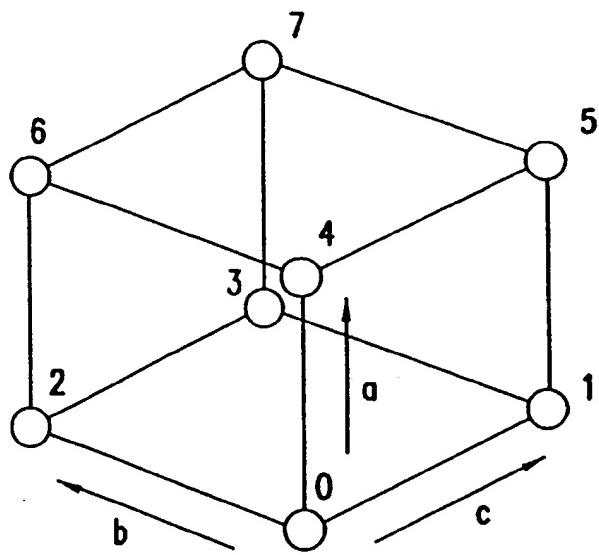


FIG.4

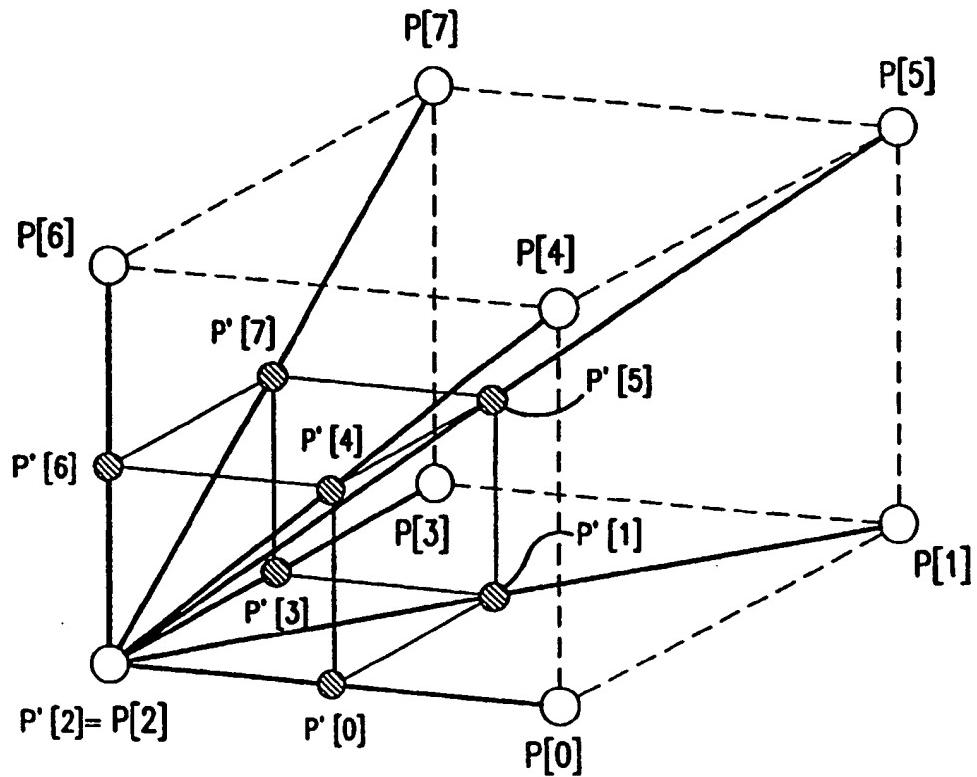


FIG.5

FIG.6B

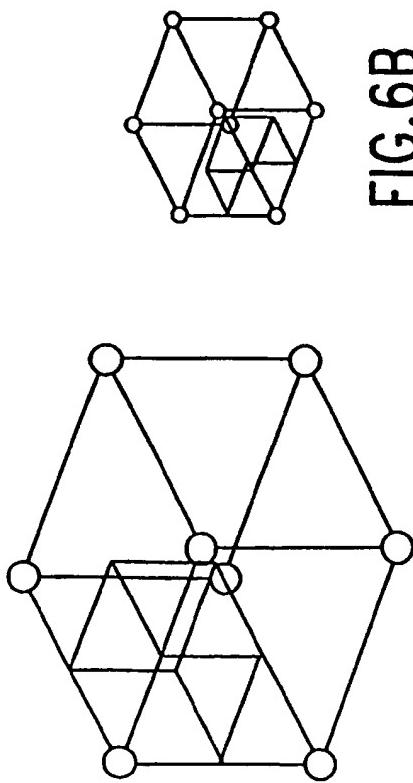
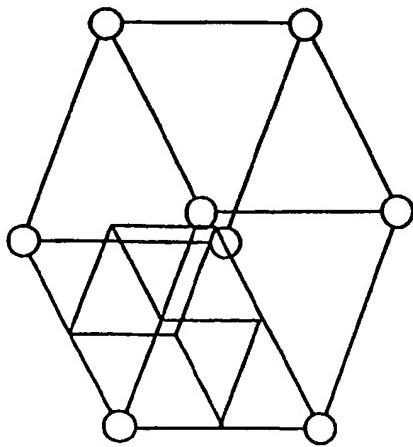


FIG.6A



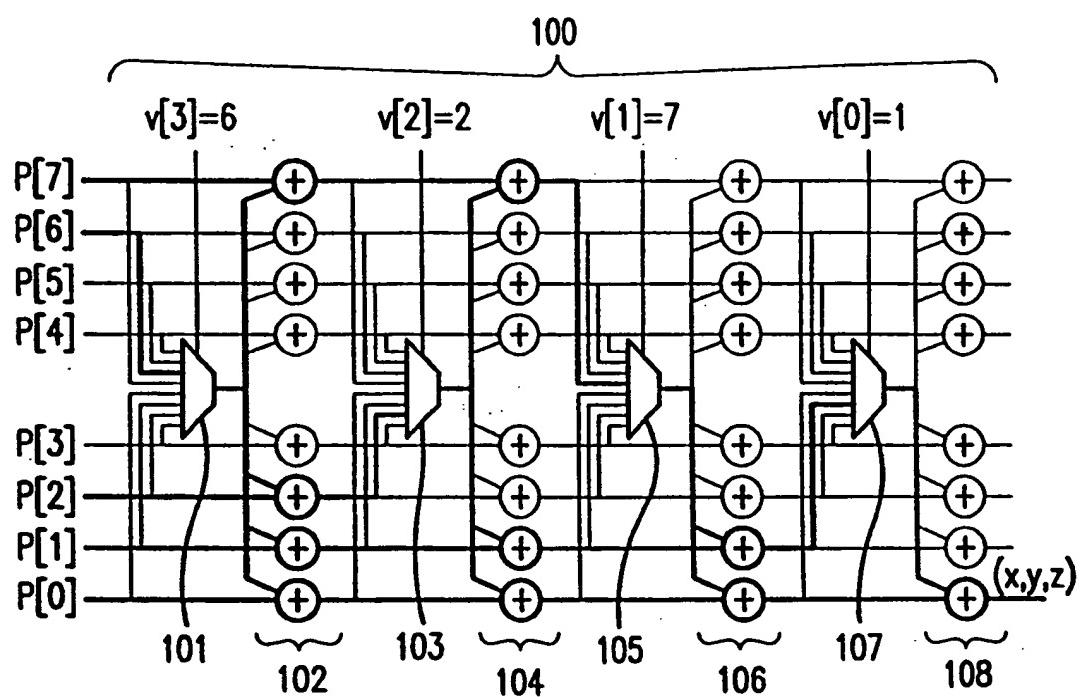


FIG.7

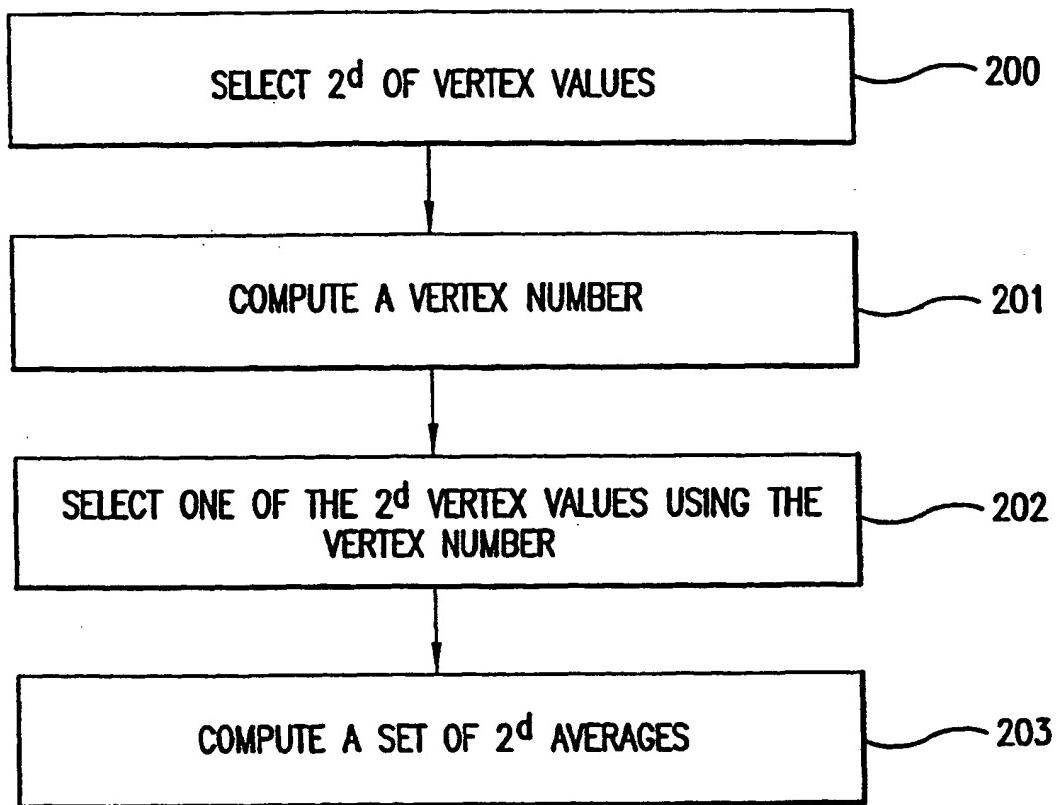
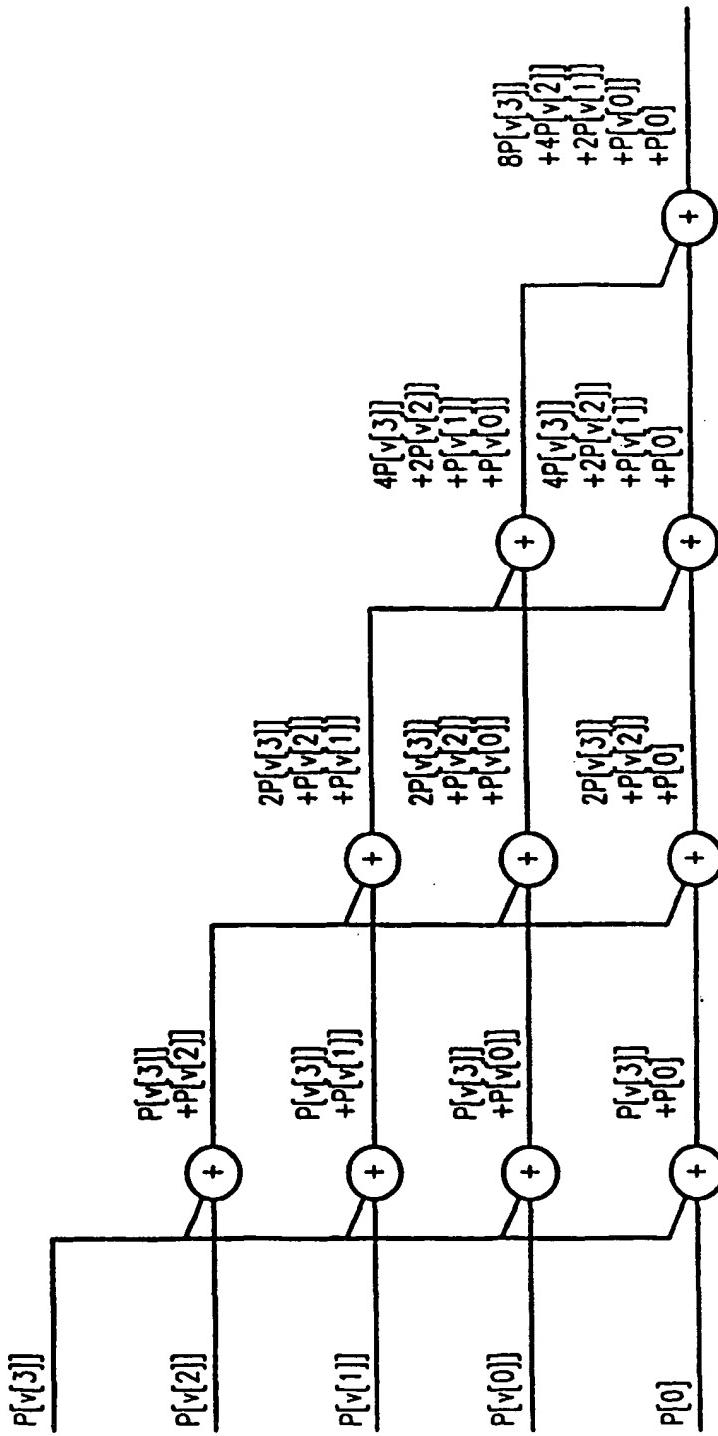


FIG.8



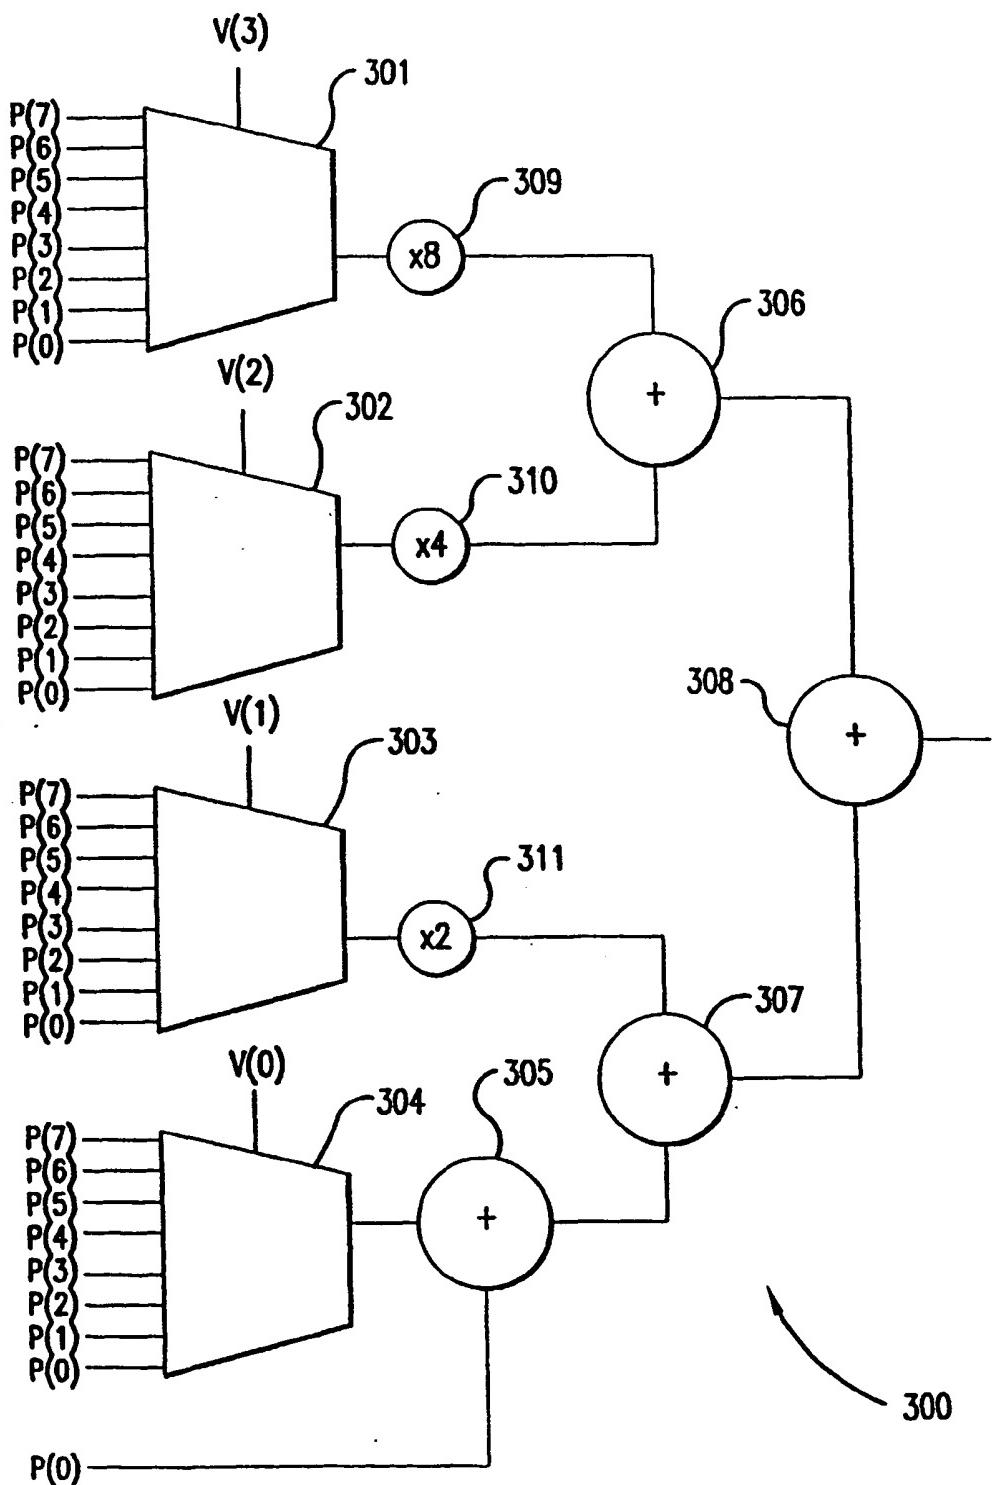


FIG.10

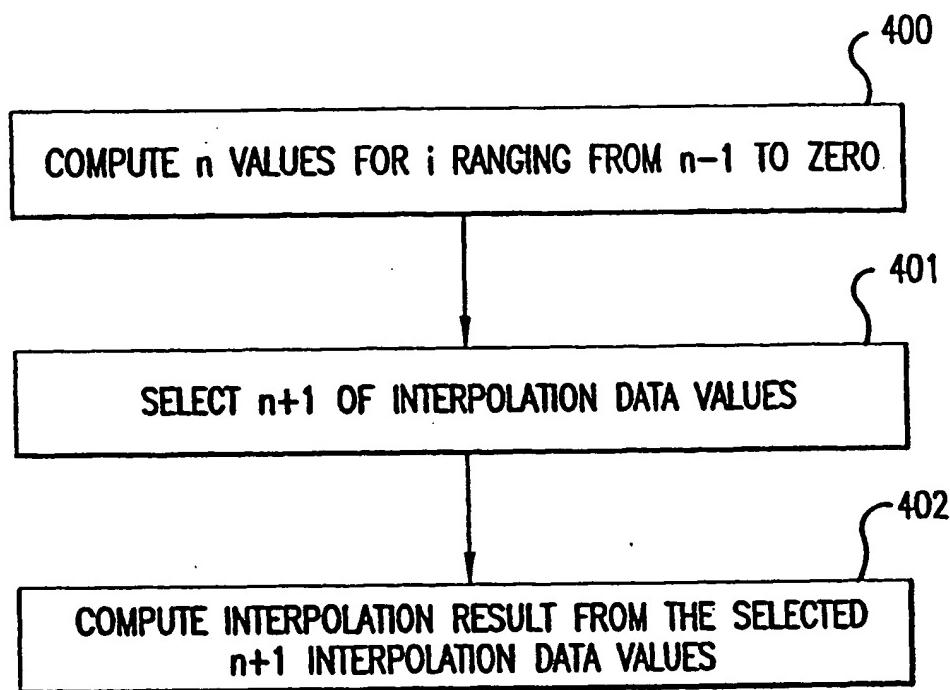


FIG.11

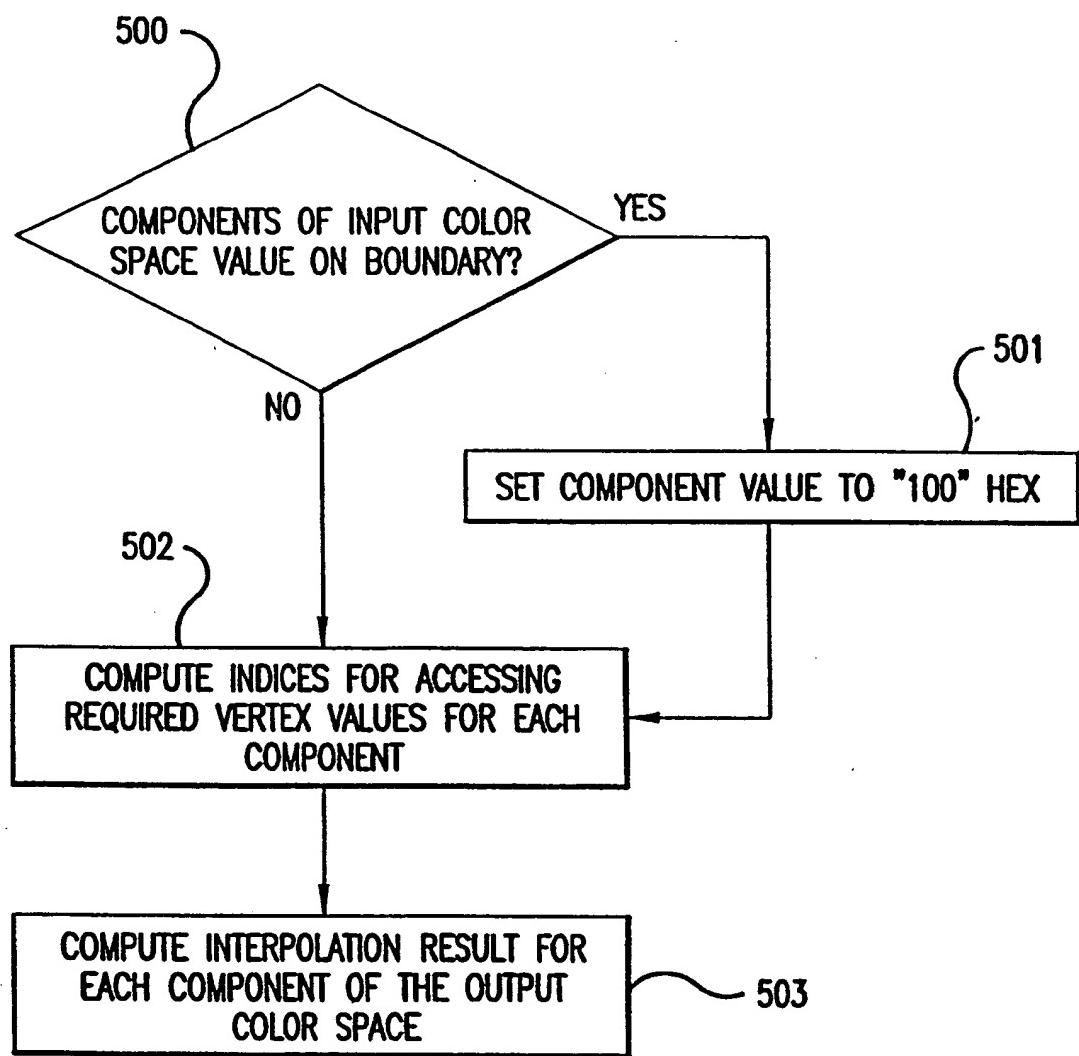


FIG.12

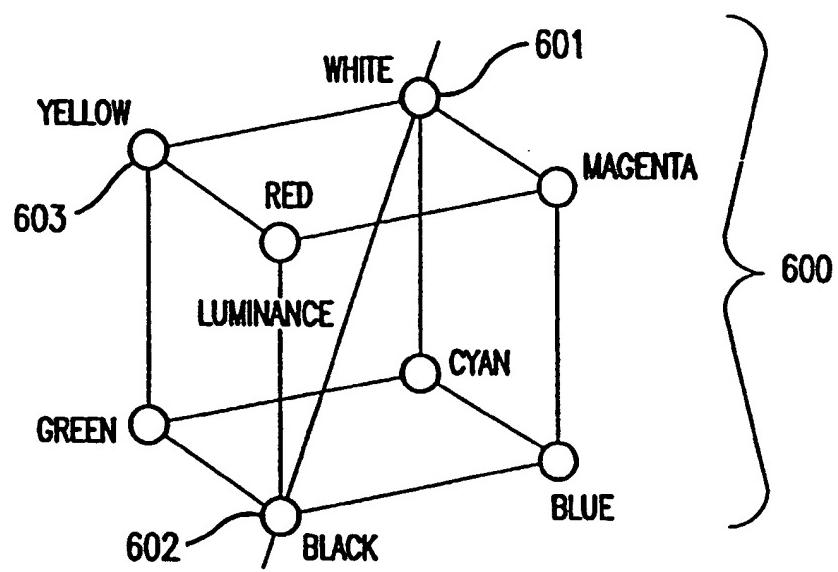


FIG.13

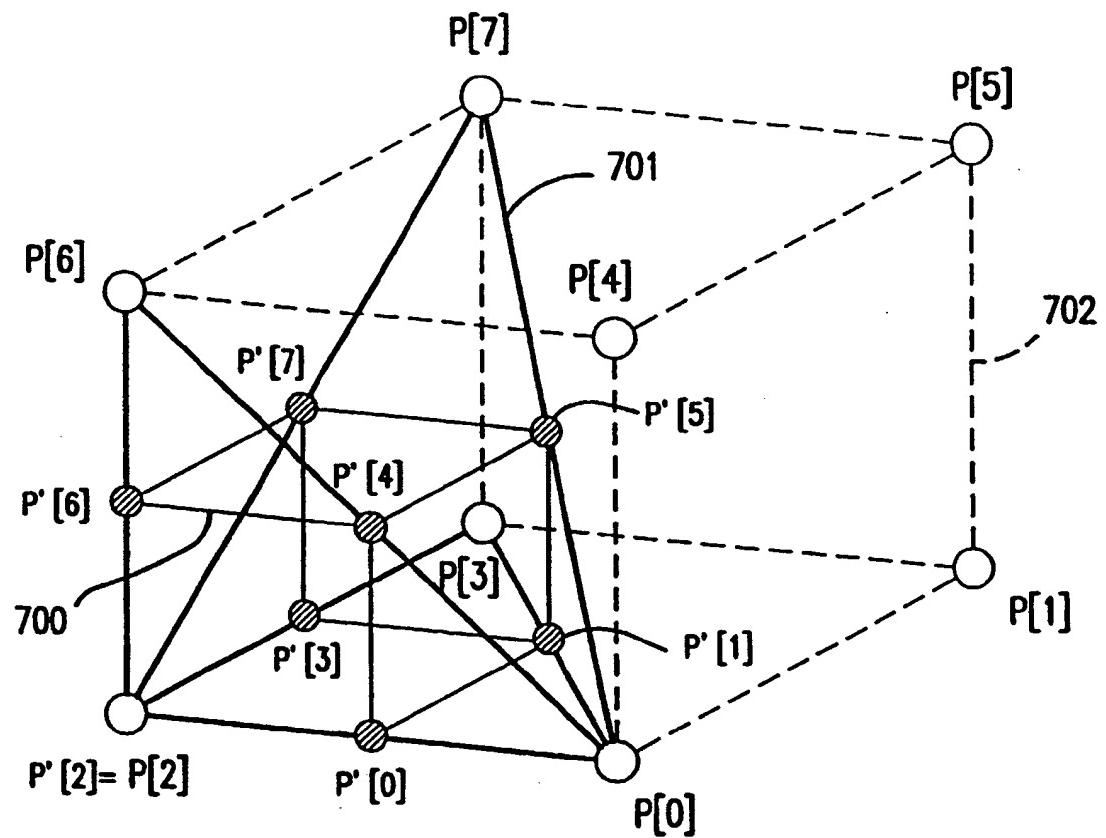


FIG. 14

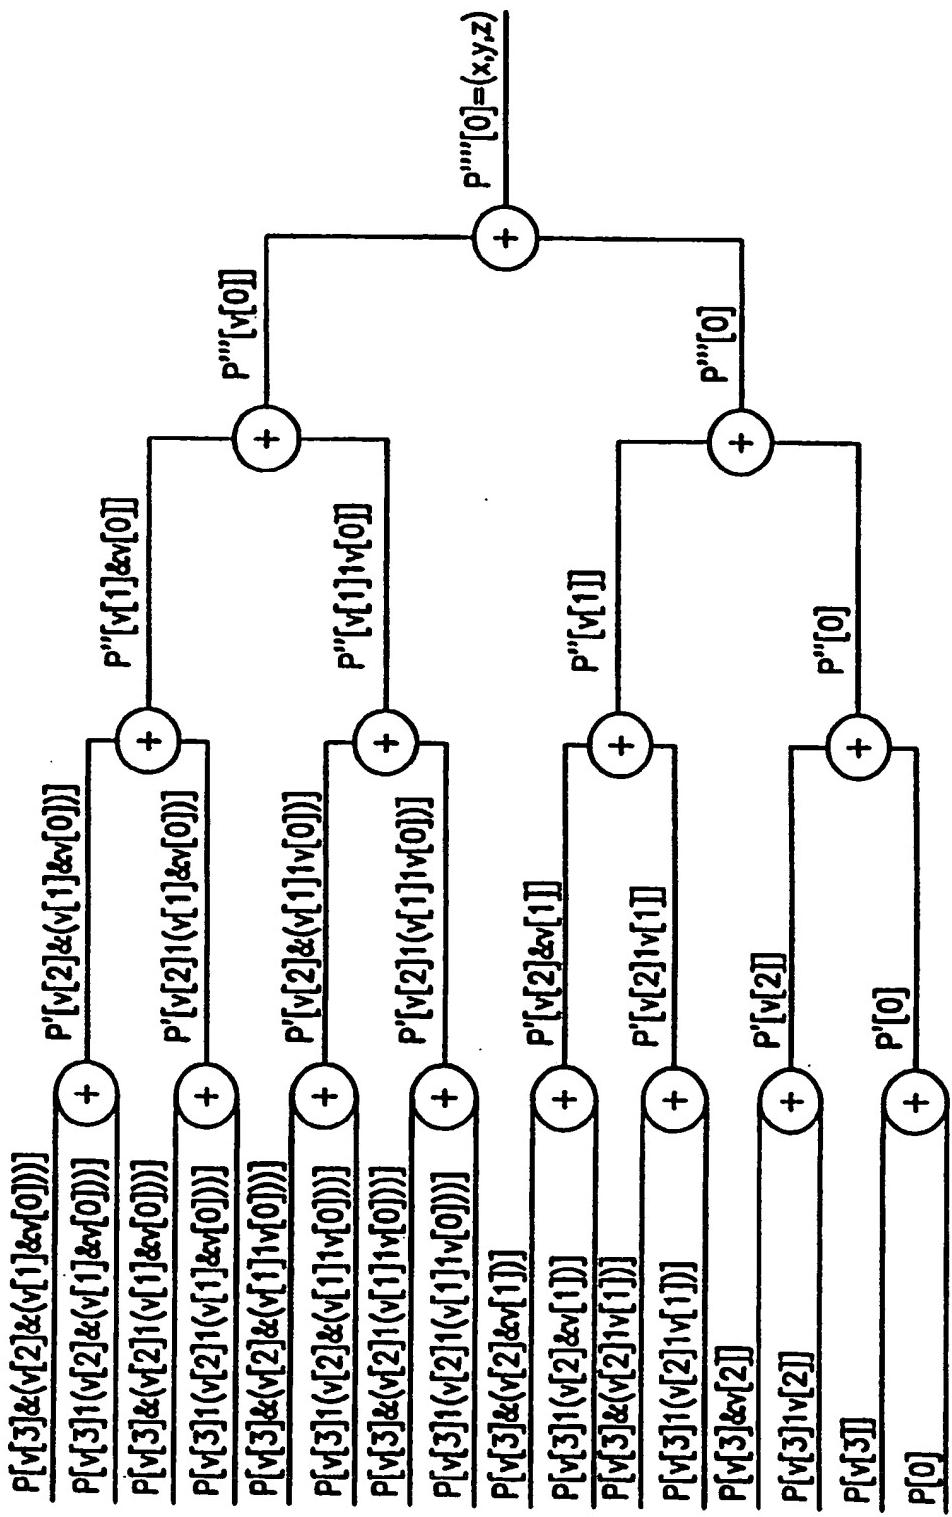


FIG. 15

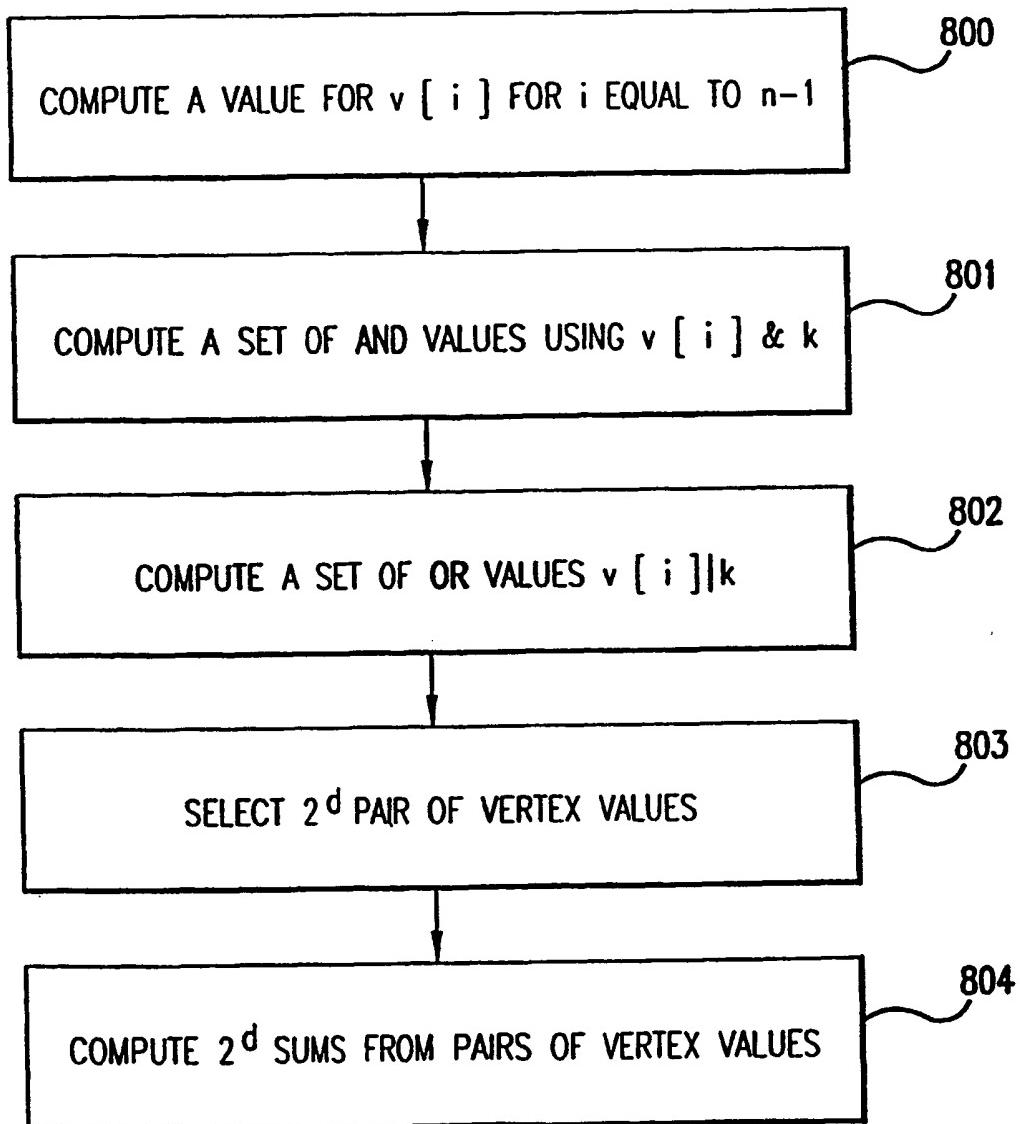


FIG.16

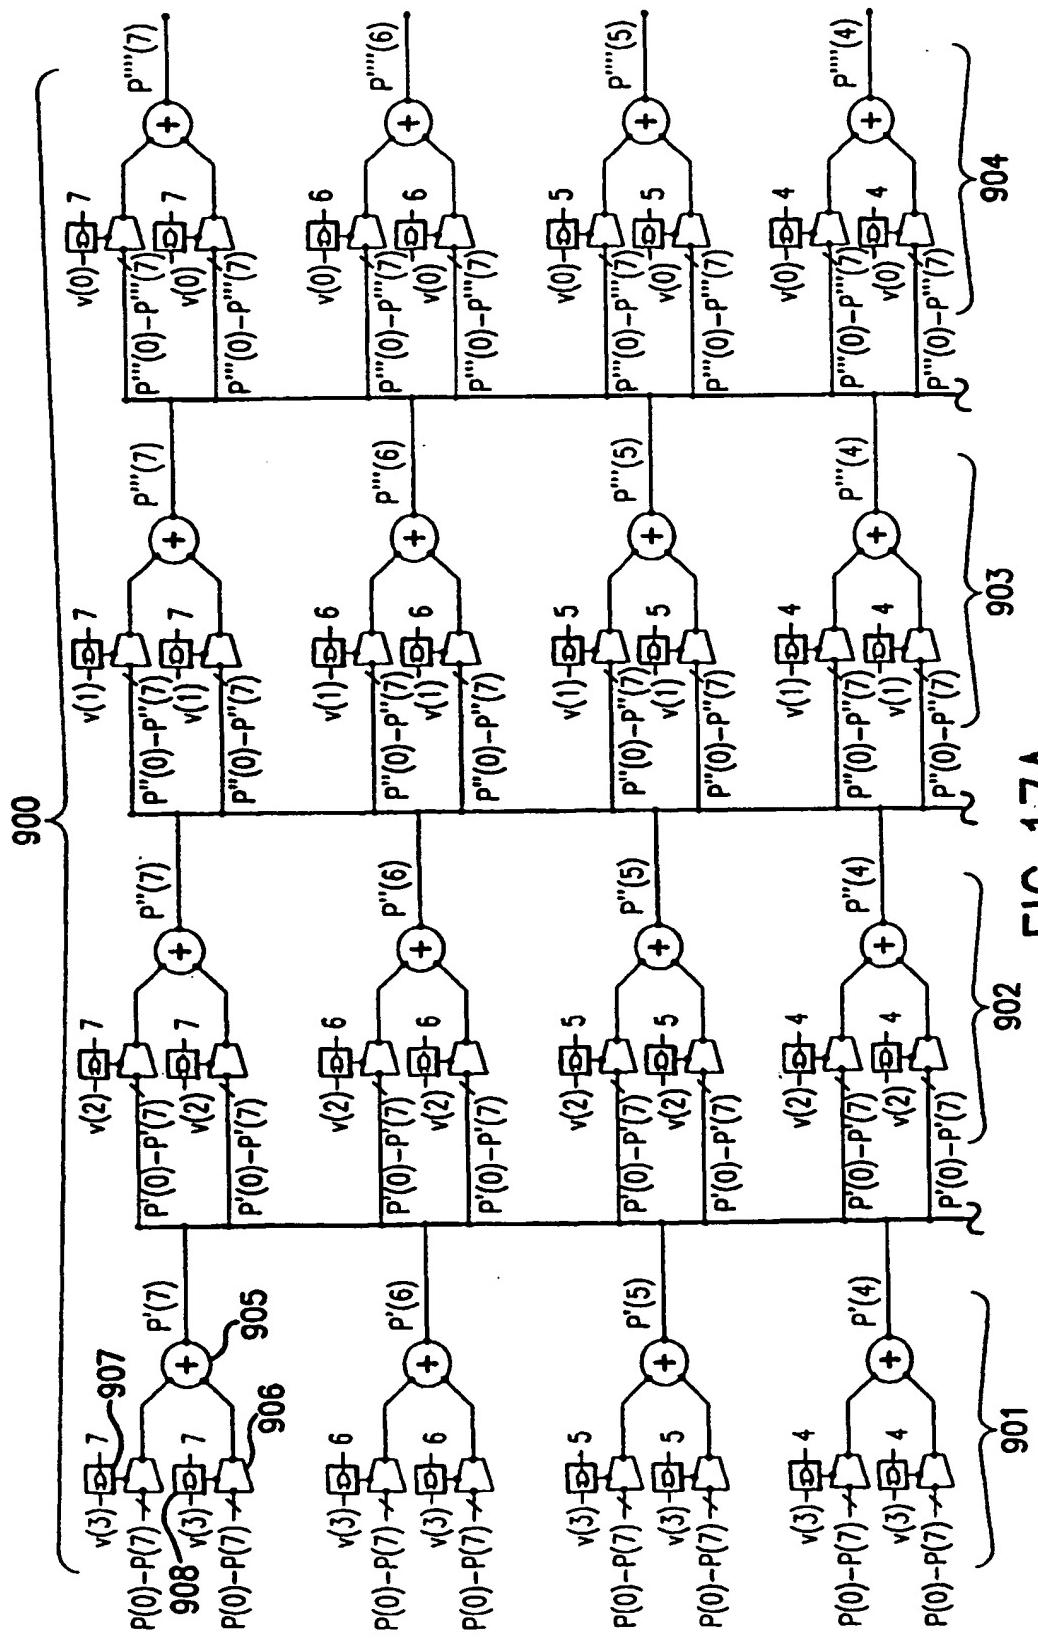


FIG. 17A

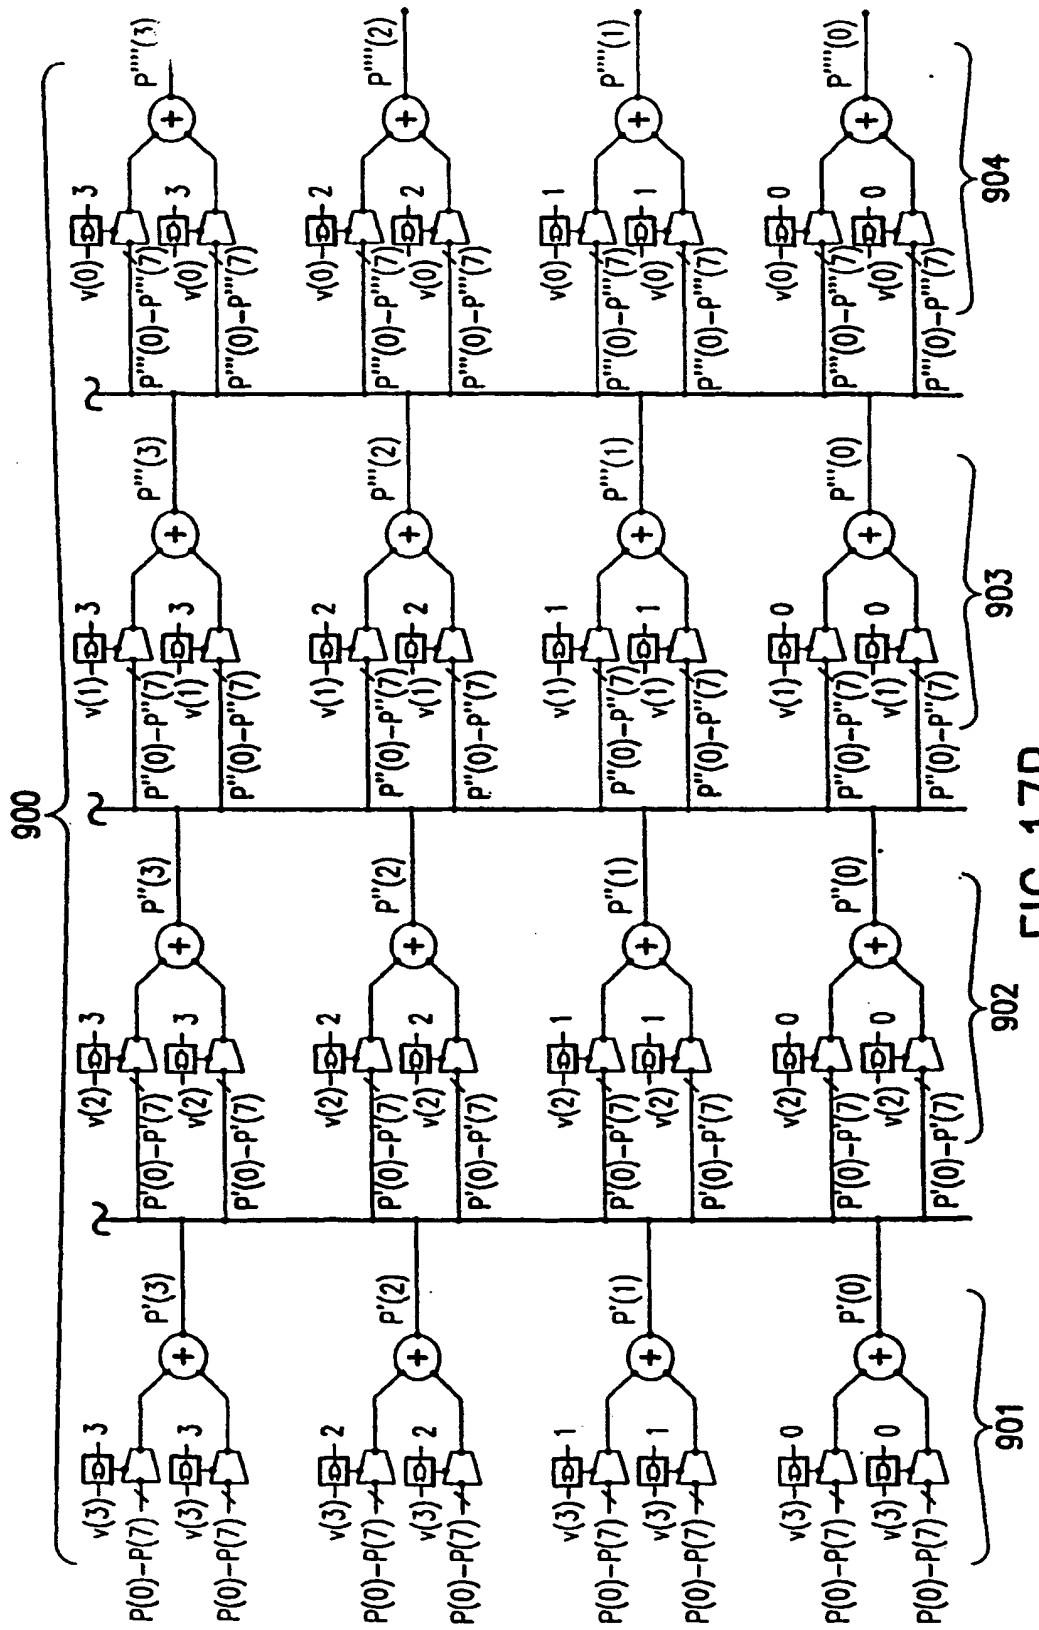
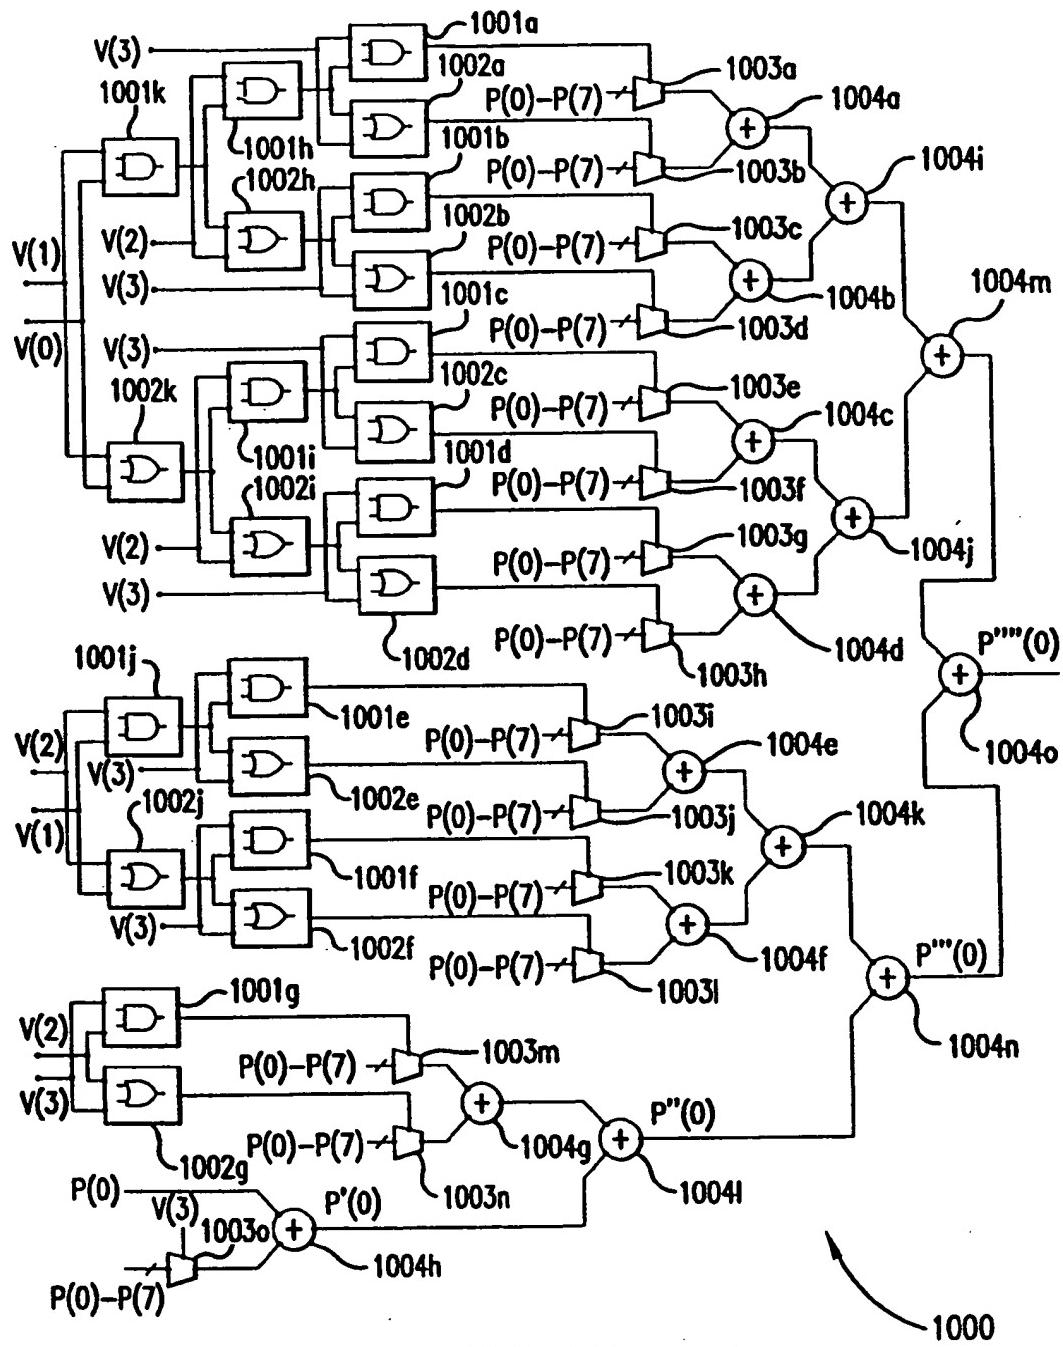


FIG. 17B



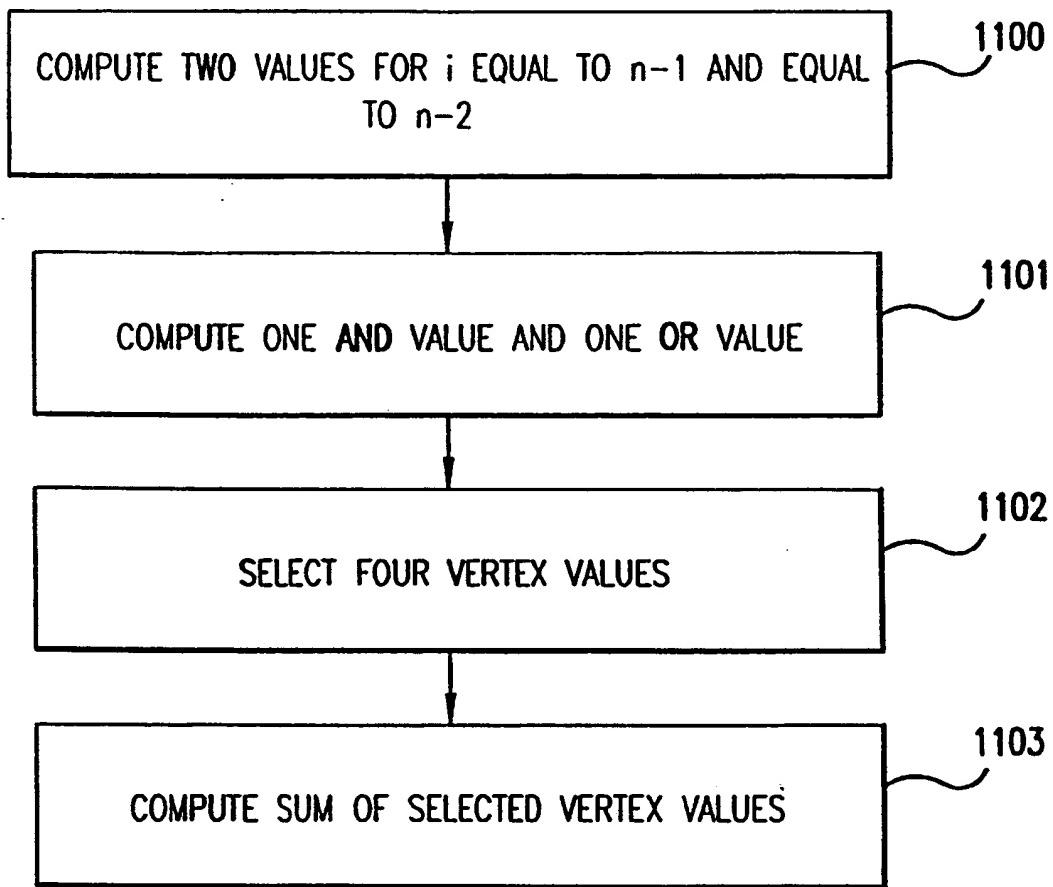


FIG.19

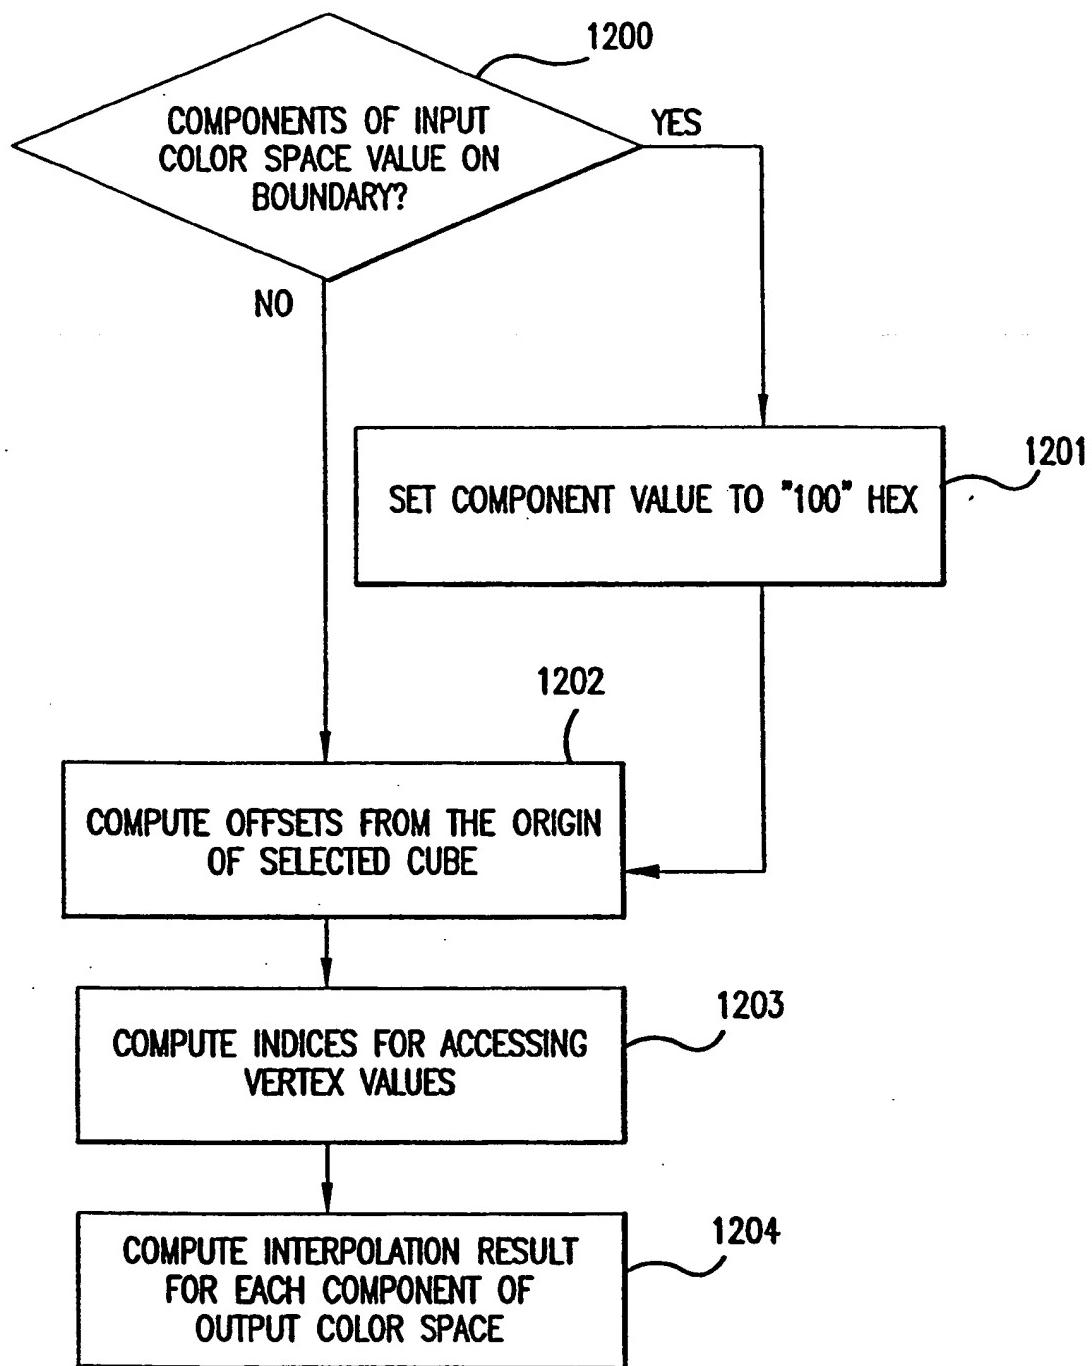


FIG.20

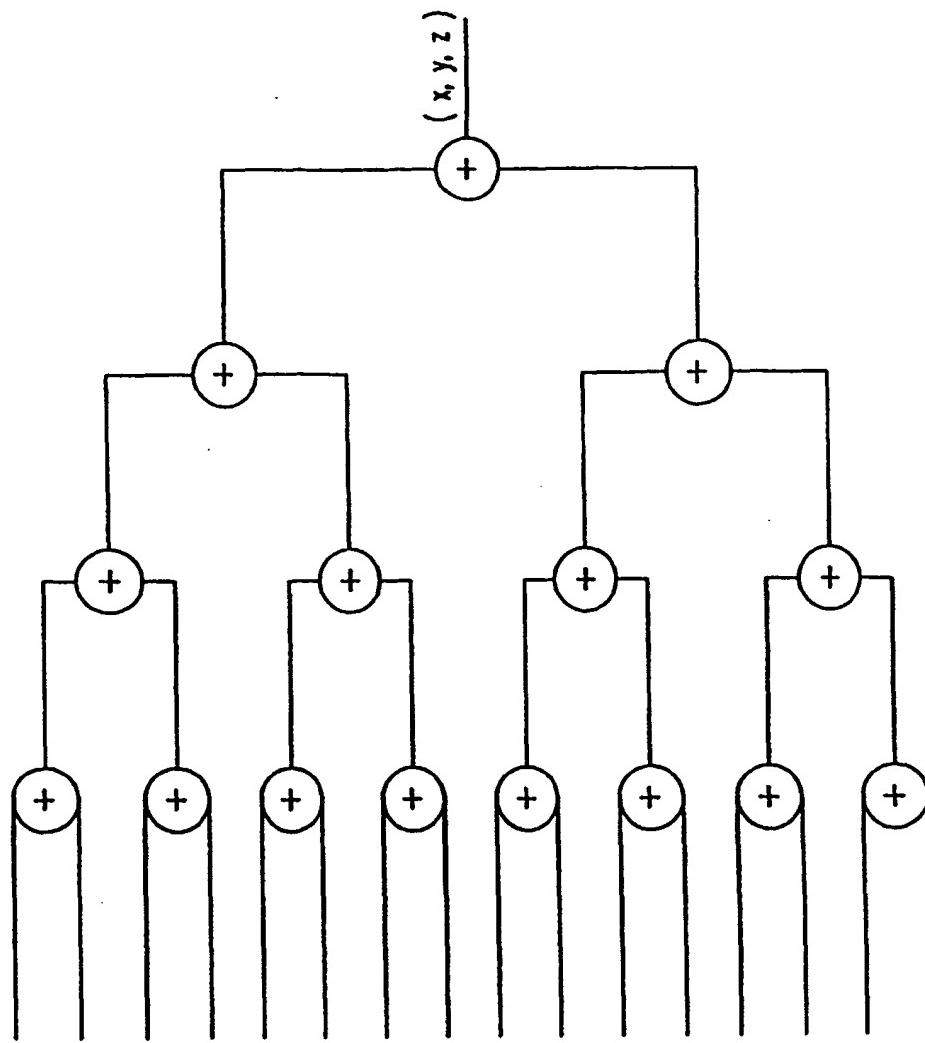


FIG.21

$P[0]$	$P[y_3] \& (y_2 \& (y_1 \& \neg y_0))$
$P[1]$	$P[y_3] \mid (y_2 \& (y_1 \& \neg y_0))$
$P[2]$	$P[y_3] \& (y_2 \mid (y_1 \& \neg y_0))$
$P[3]$	$P[y_3] \mid (y_2 \mid (y_1 \& \neg y_0))$
$P[4]$	$P[y_3] \& (y_2 \& (y_1 \mid \neg y_0))$
$P[5]$	$P[y_3] \mid (y_2 \& (y_1 \mid \neg y_0))$
$P[6]$	$P[y_3] \& (y_2 \mid (y_1 \mid \neg y_0))$
$P[7]$	$P[y_3] \mid (y_2 \mid (y_1 \mid \neg y_0))$
$P[8]$	$P[y_3] \& (y_2 \& \neg y_1)$
$P[9]$	$P[y_3] \mid (y_2 \& \neg y_1)$
$P[10]$	$P[y_3] \& (\neg y_2 \& \neg y_1)$
$P[11]$	$P[y_3] \mid (\neg y_2 \& \neg y_1)$
$P[12]$	$P[y_3] \& \neg y_2$
$P[13]$	$P[y_3] \mid \neg y_2$
$P[14]$	$P[0]$

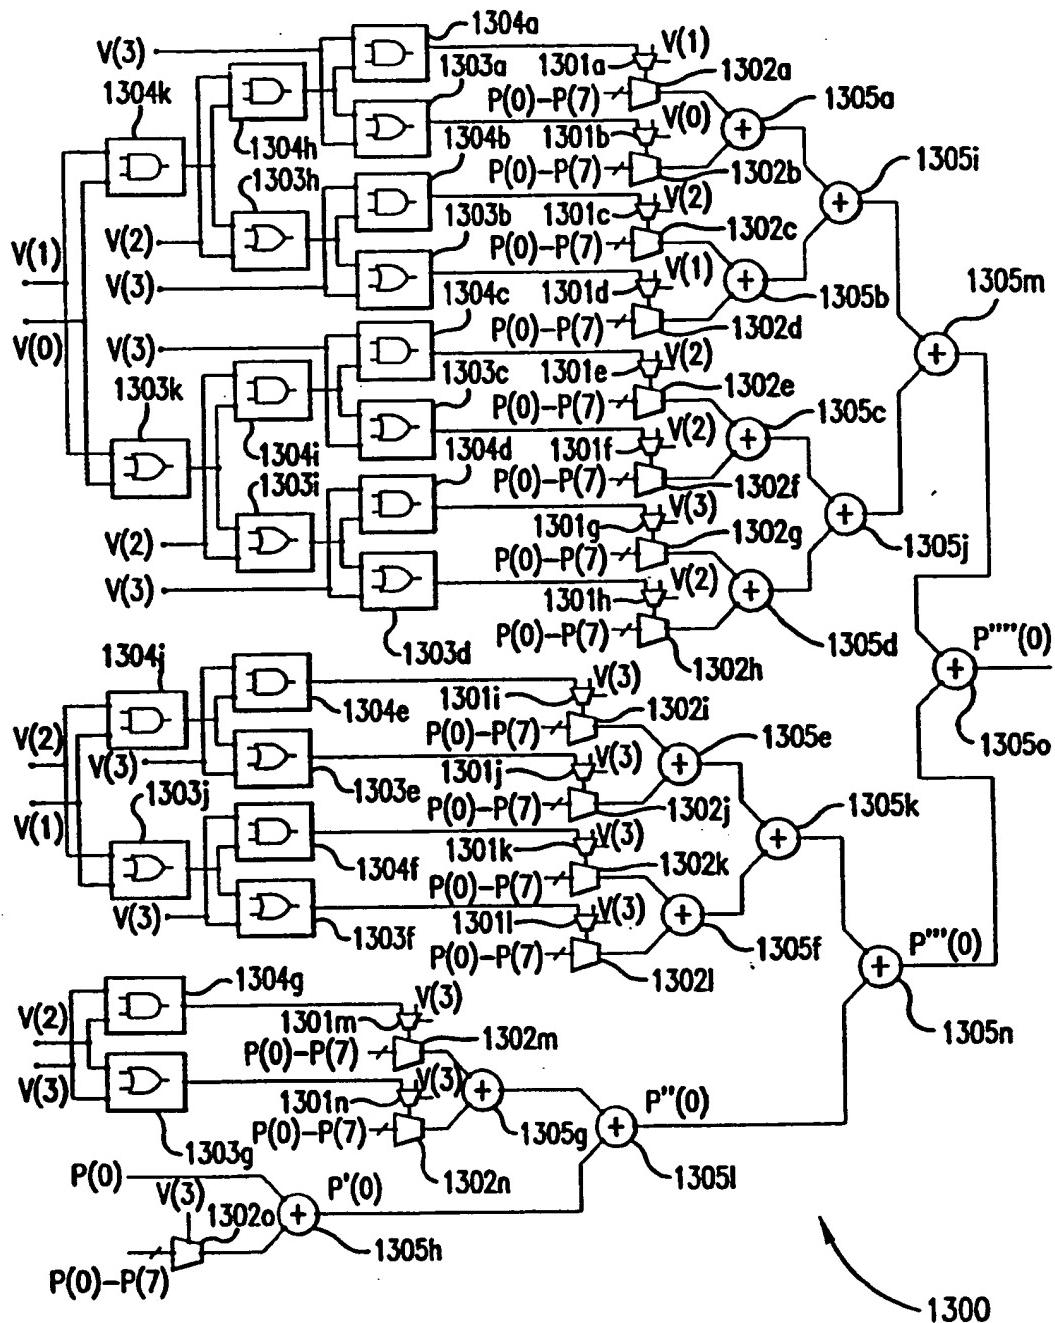


FIG. 22

FIG.23E



FIG.23D

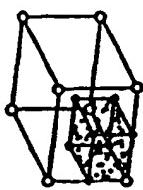


FIG.23C

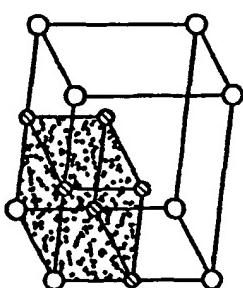


FIG.23B

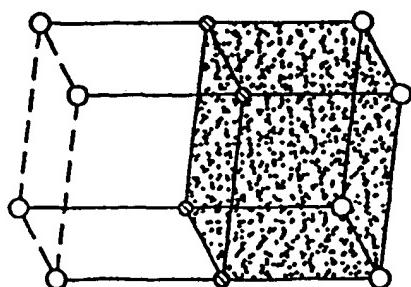


FIG.23A

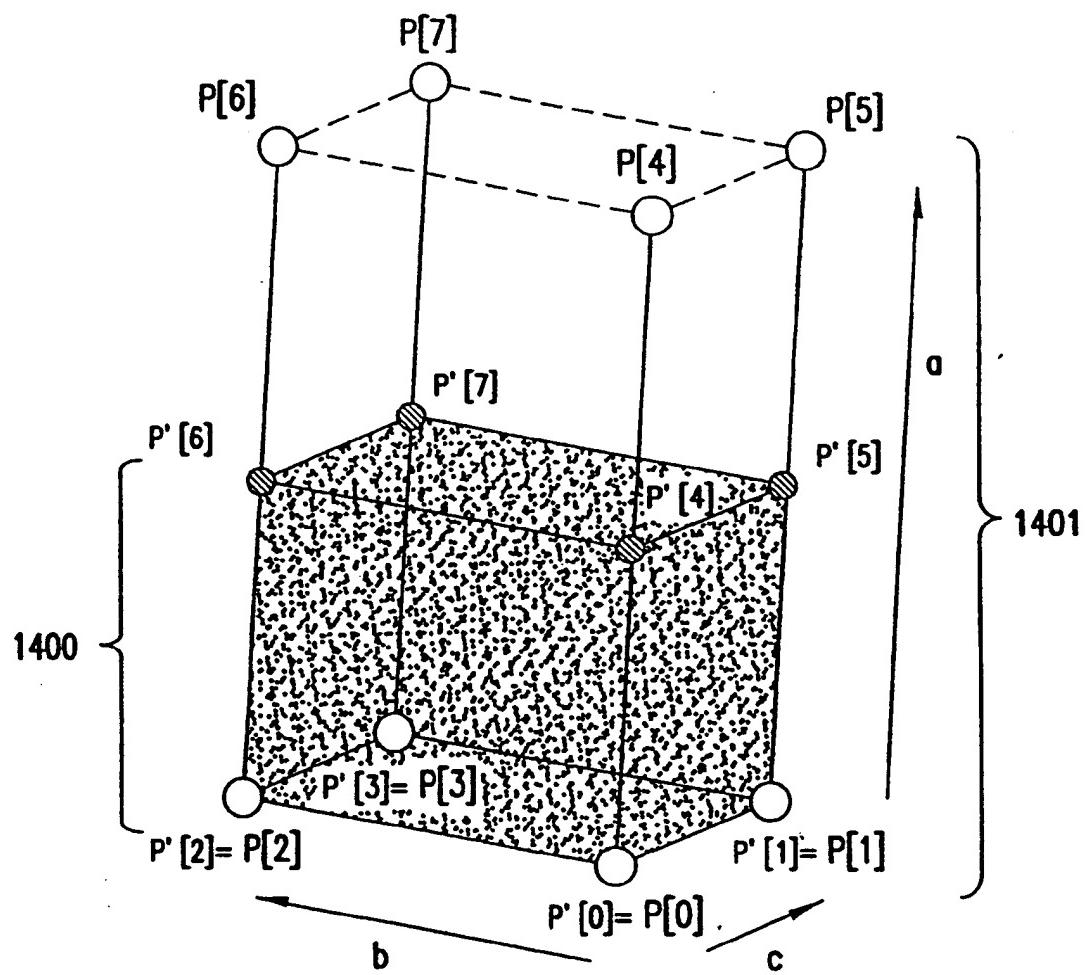


FIG.24

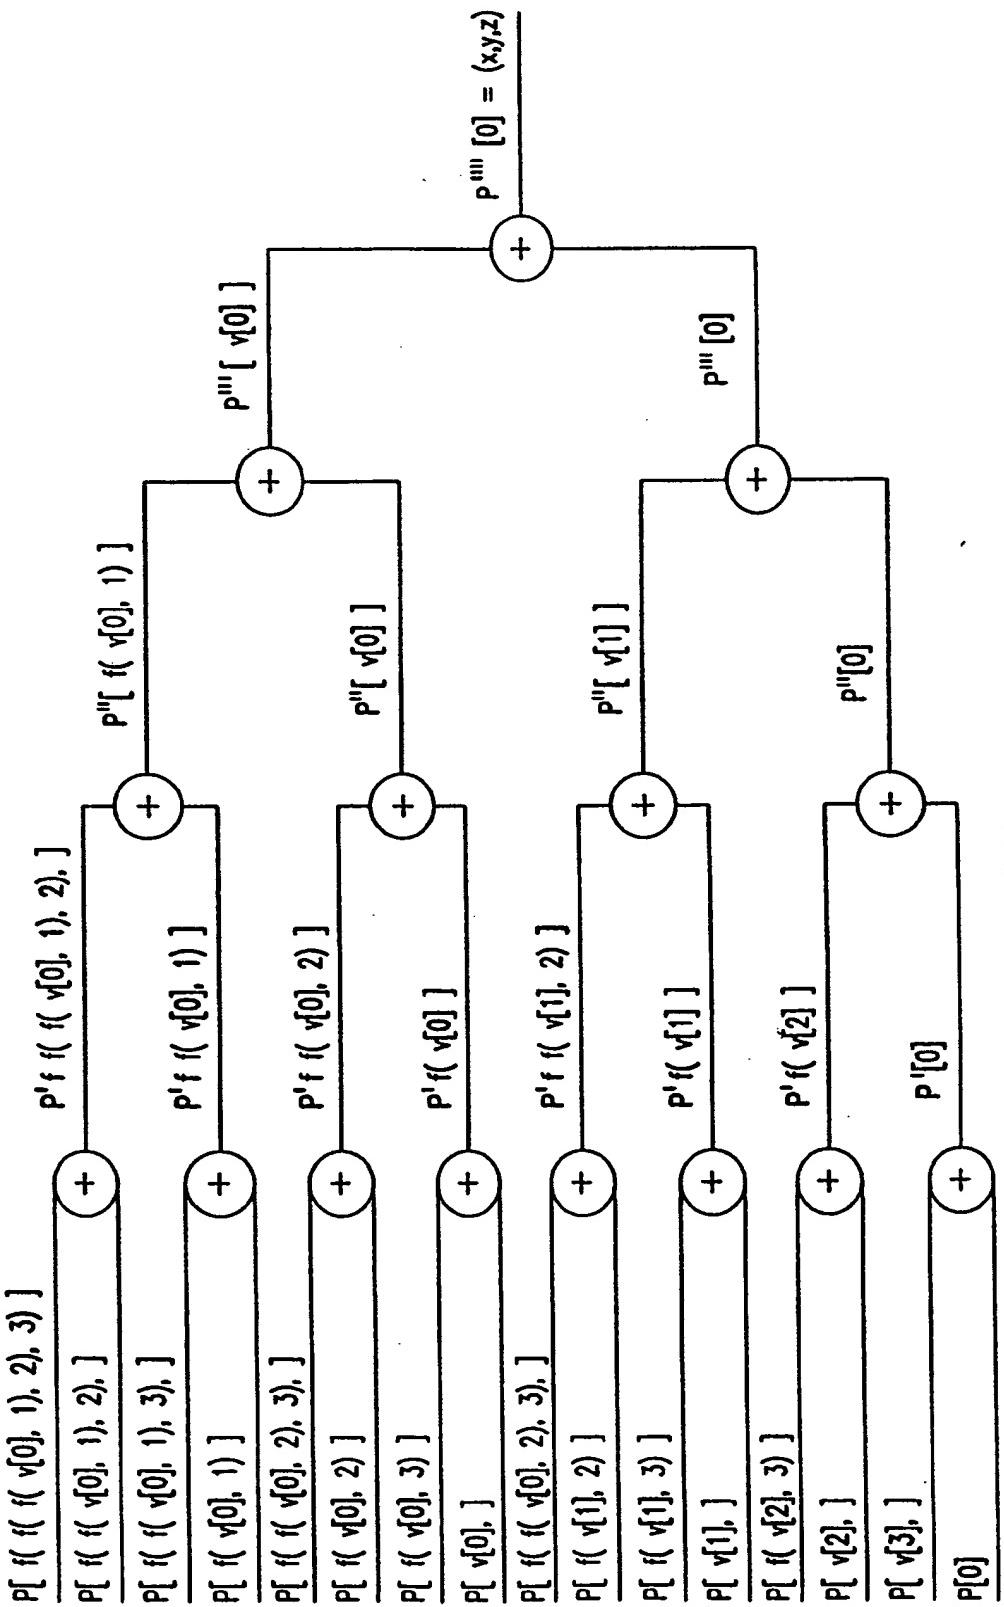


FIG. 25

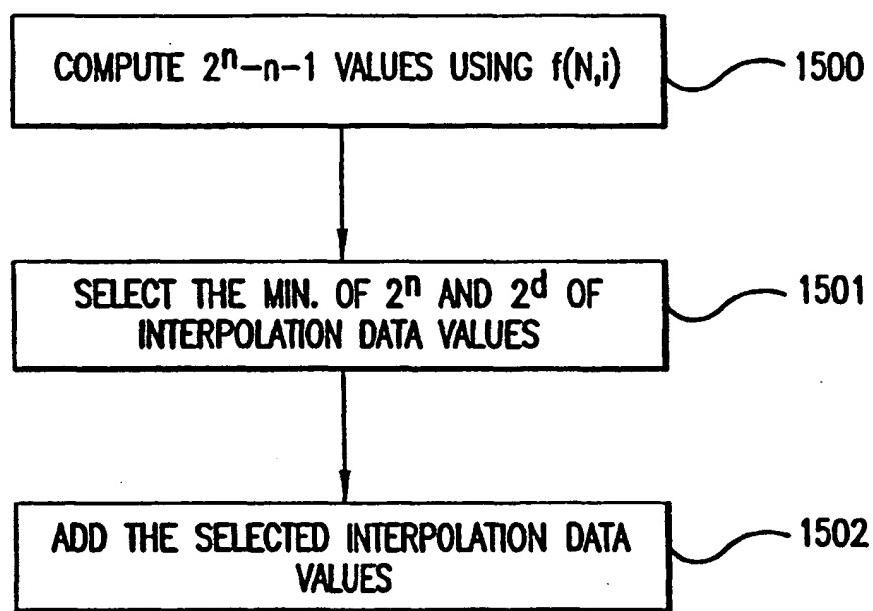


FIG. 26

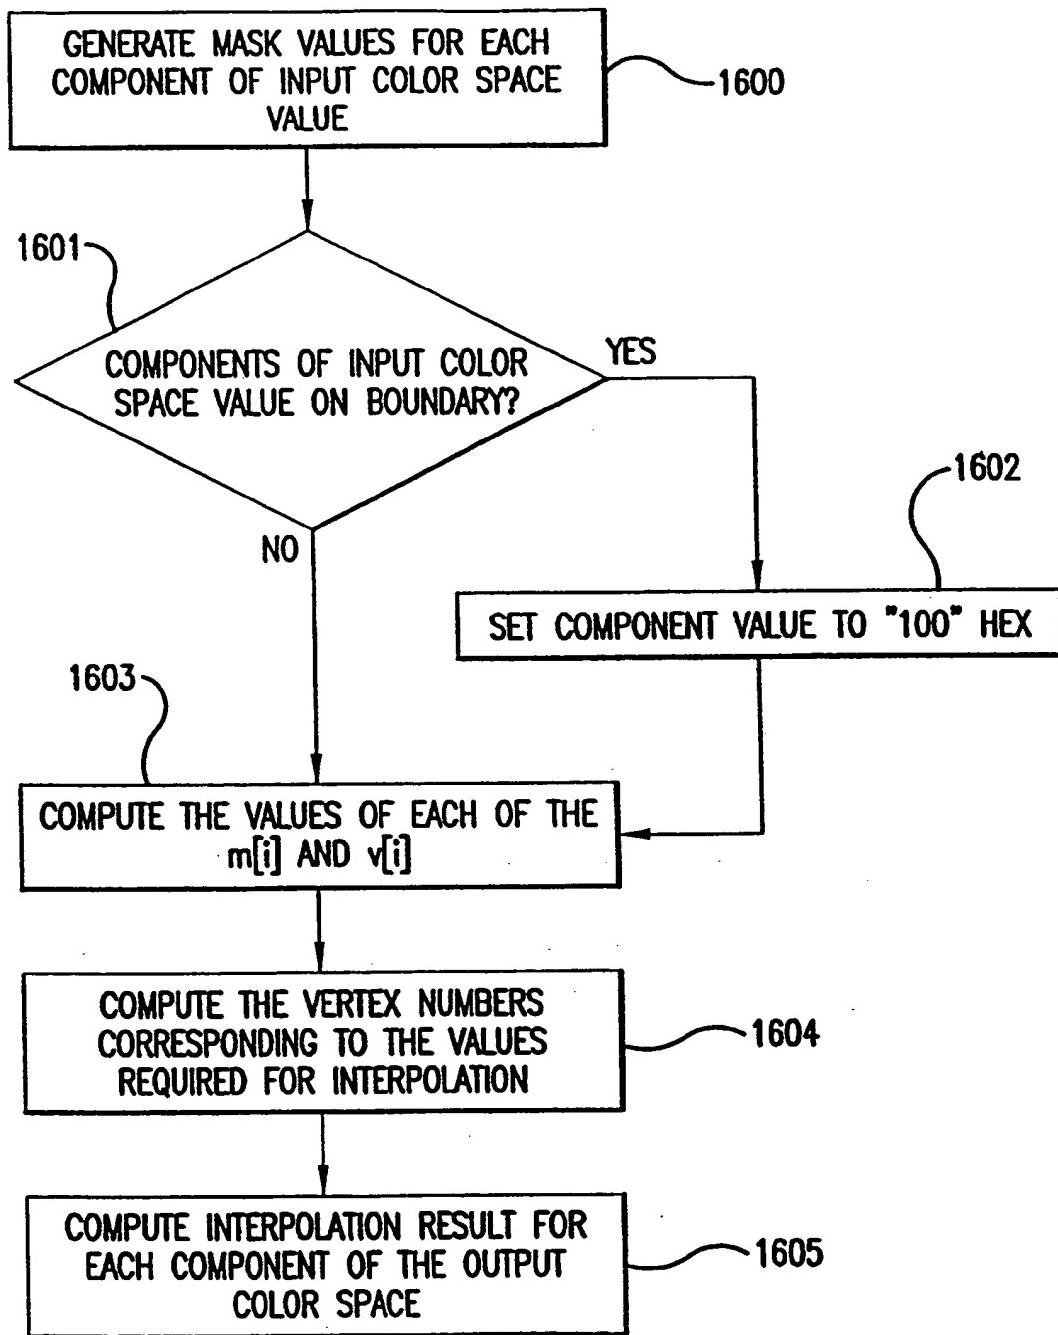
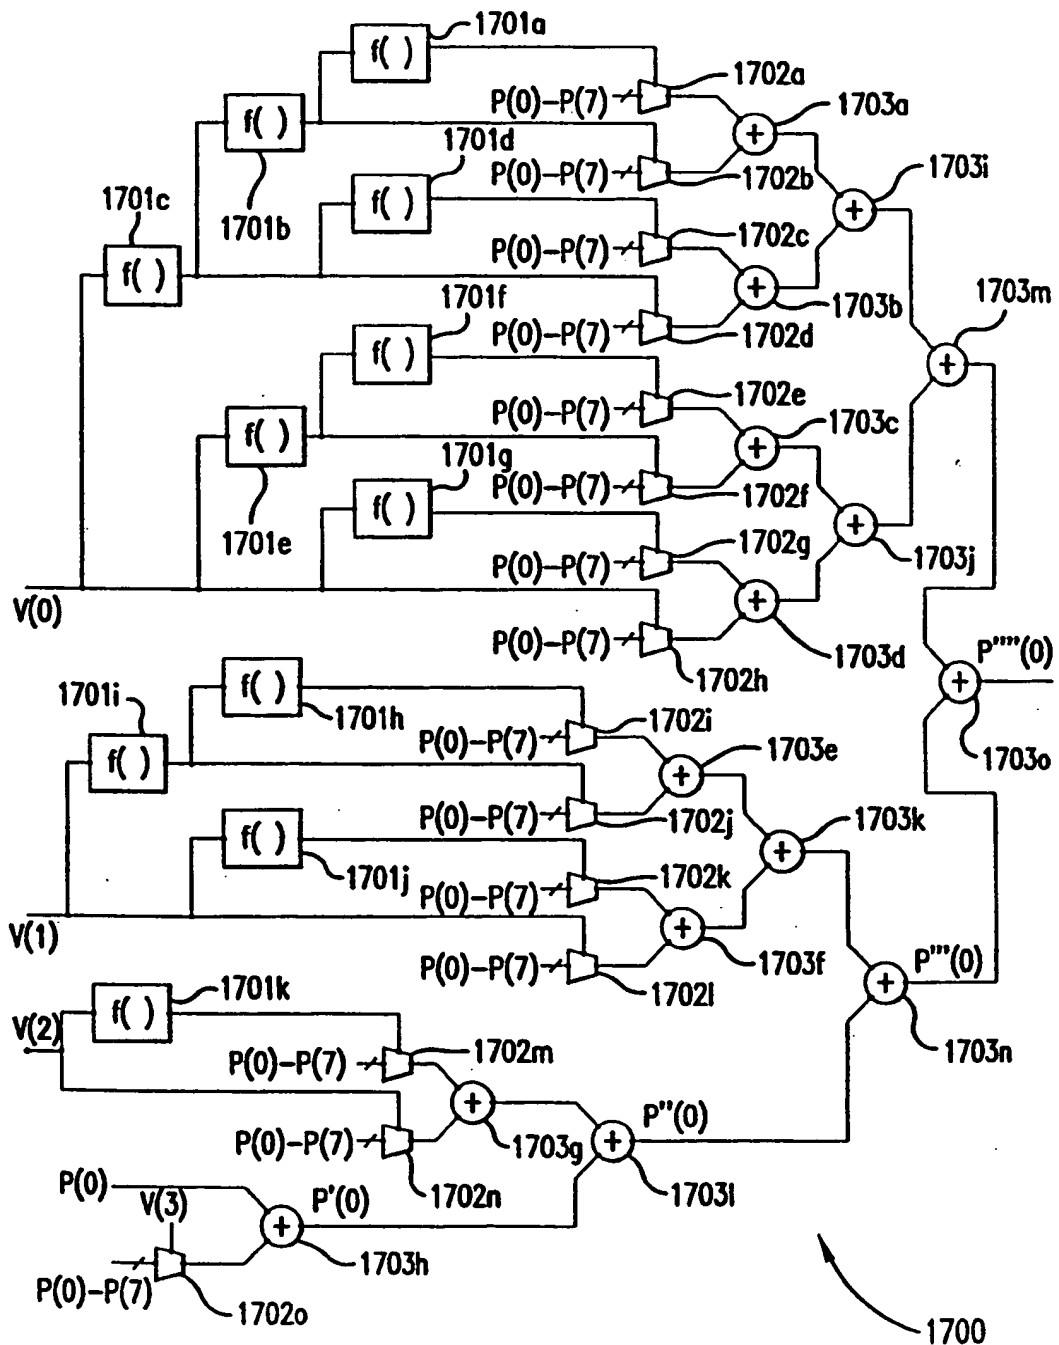


FIG. 27



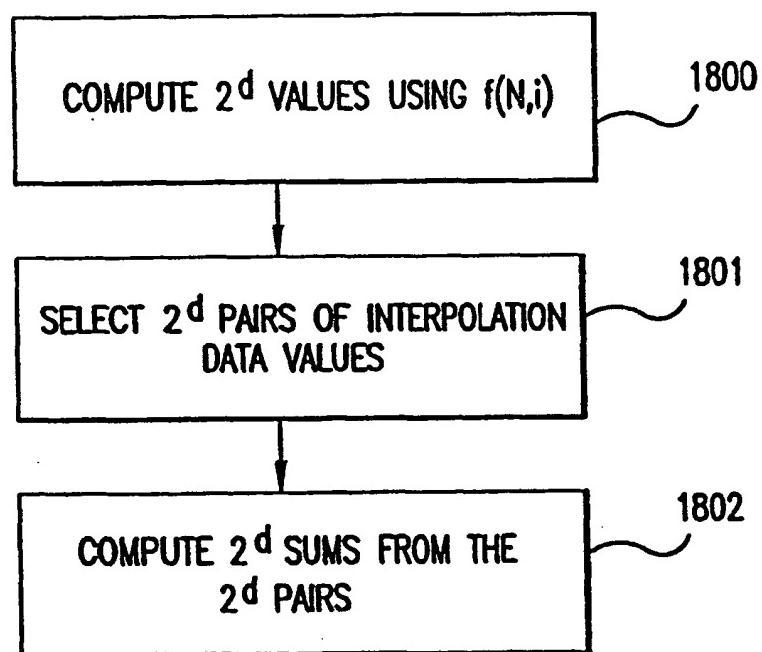


FIG.29

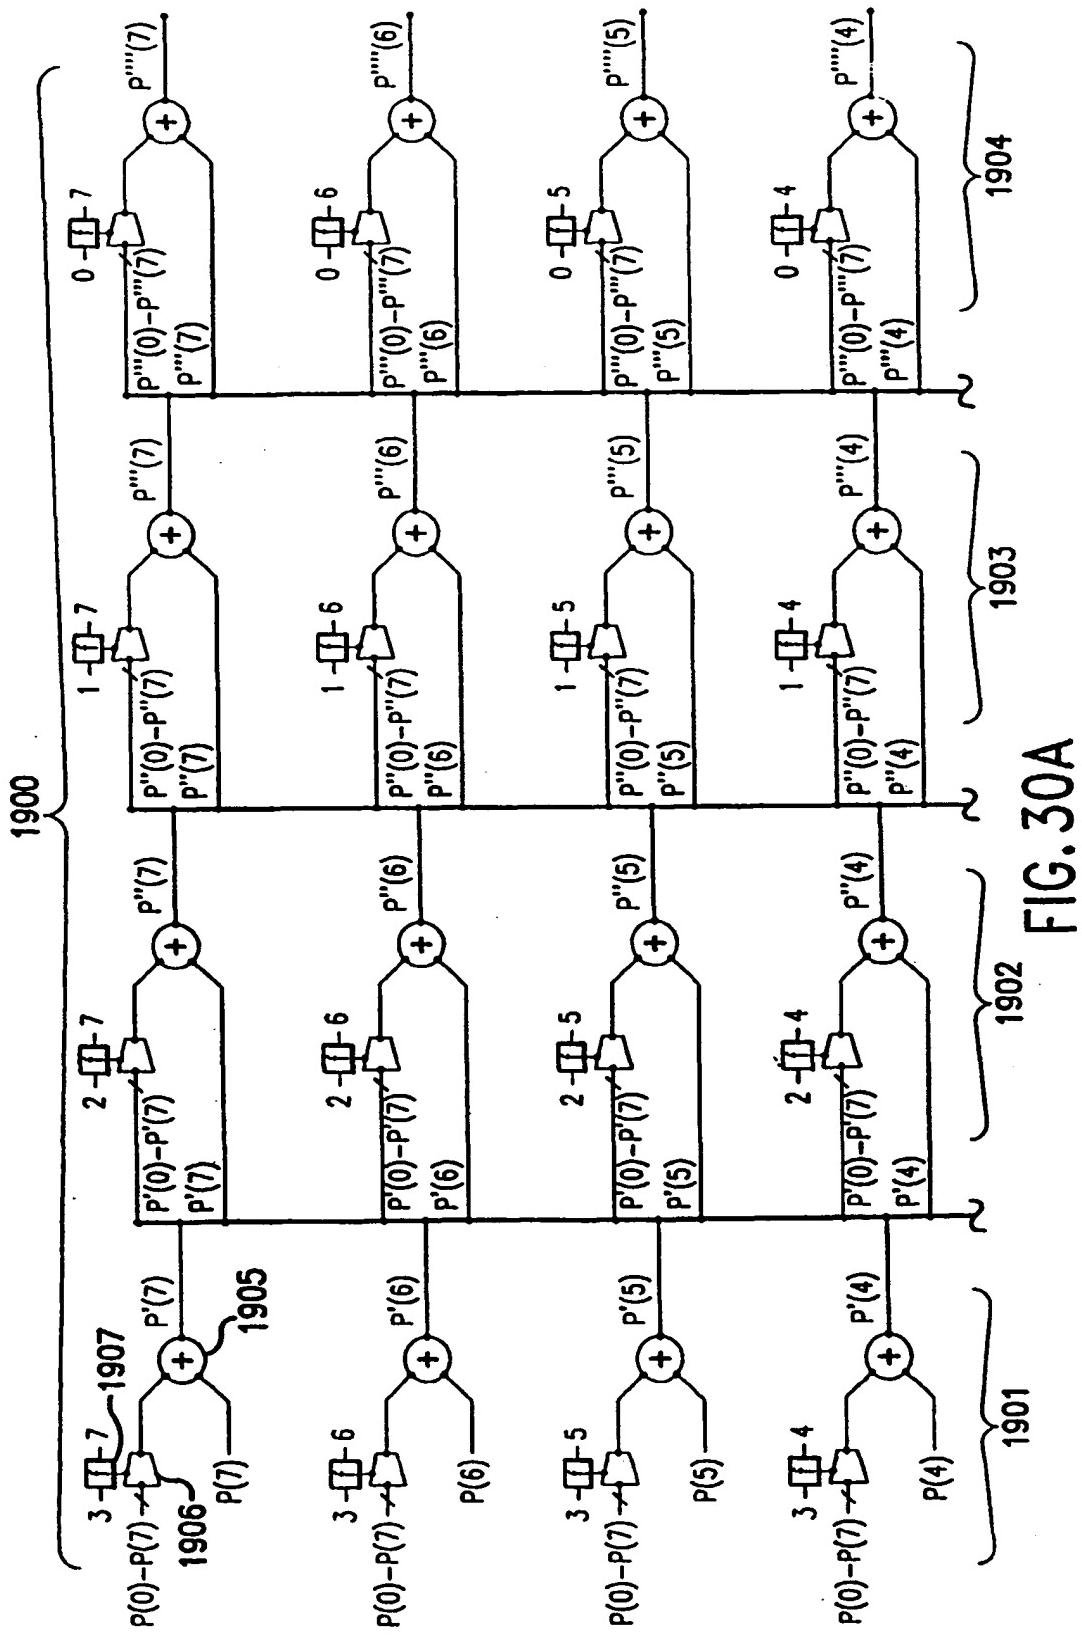
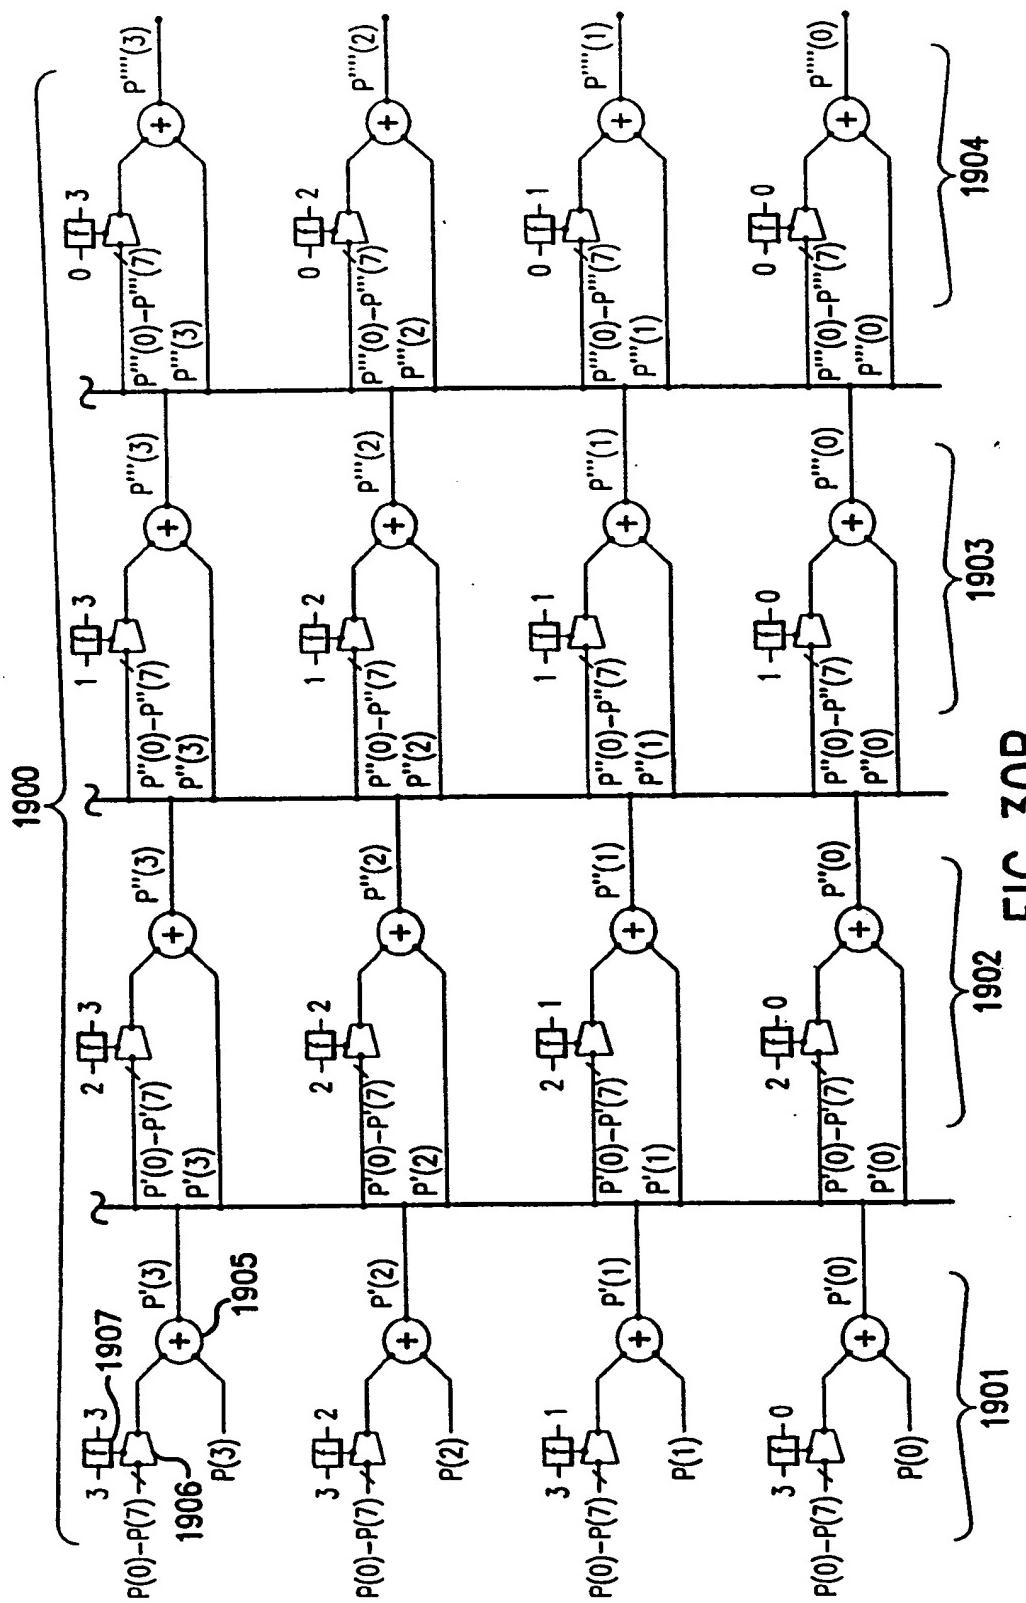


FIG. 30A



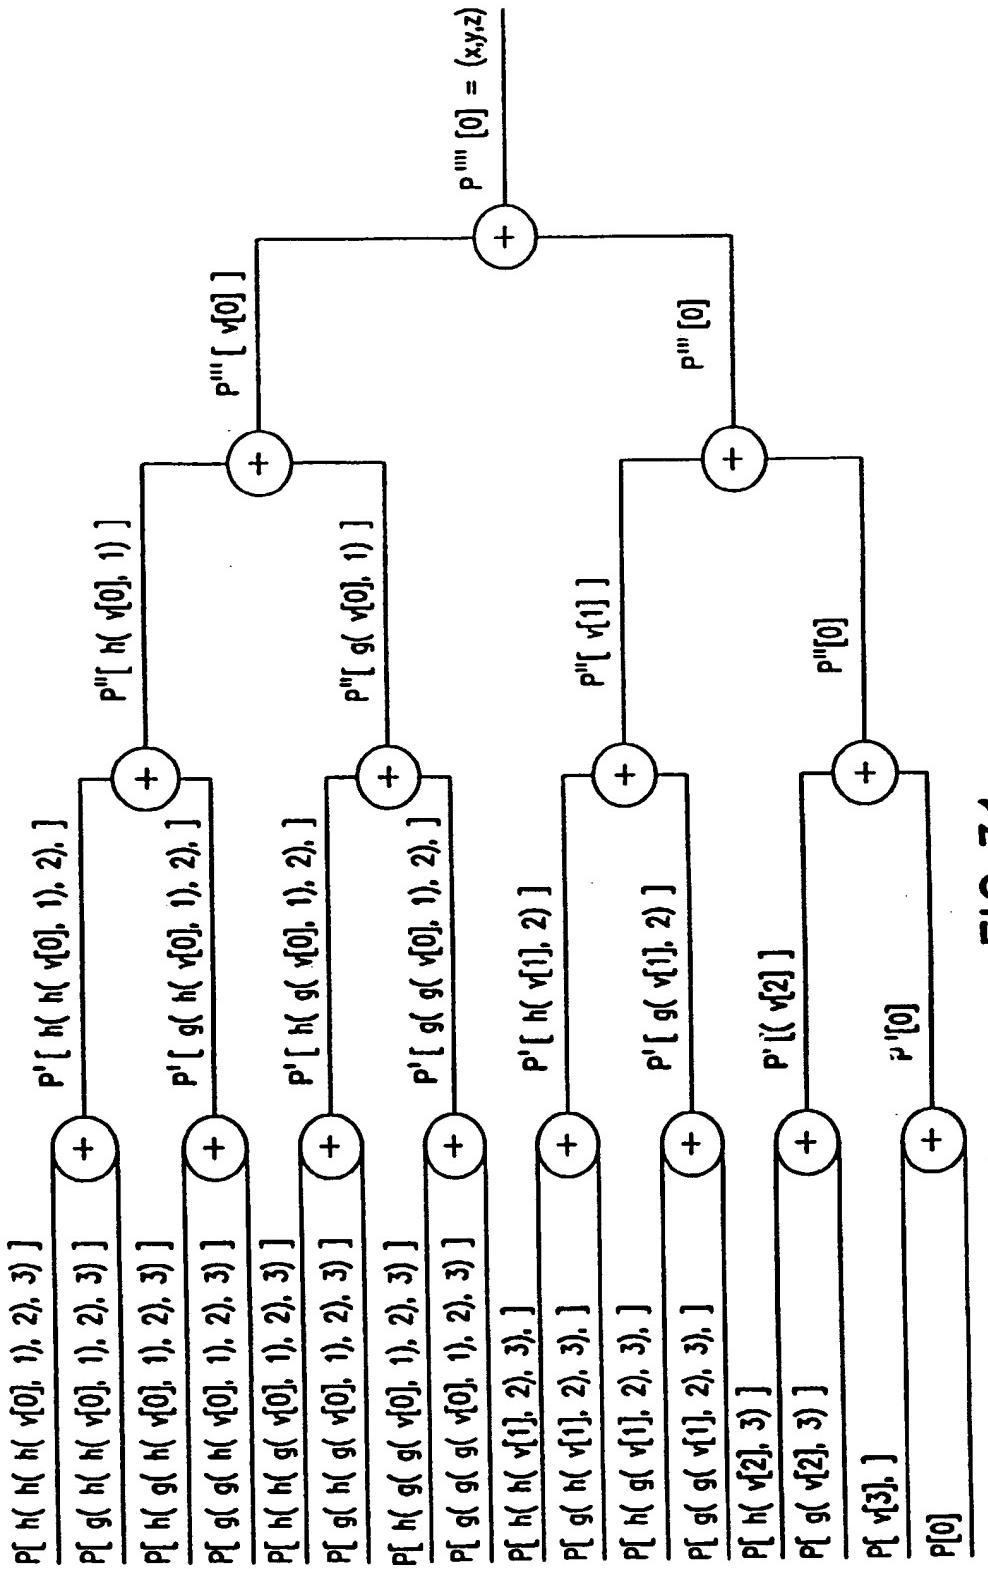


FIG. 31

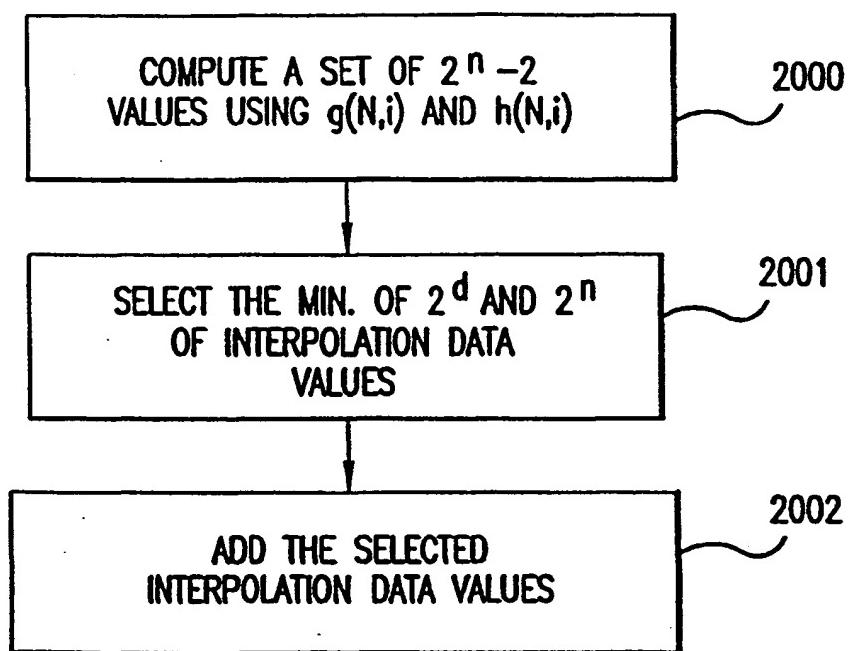


FIG.32

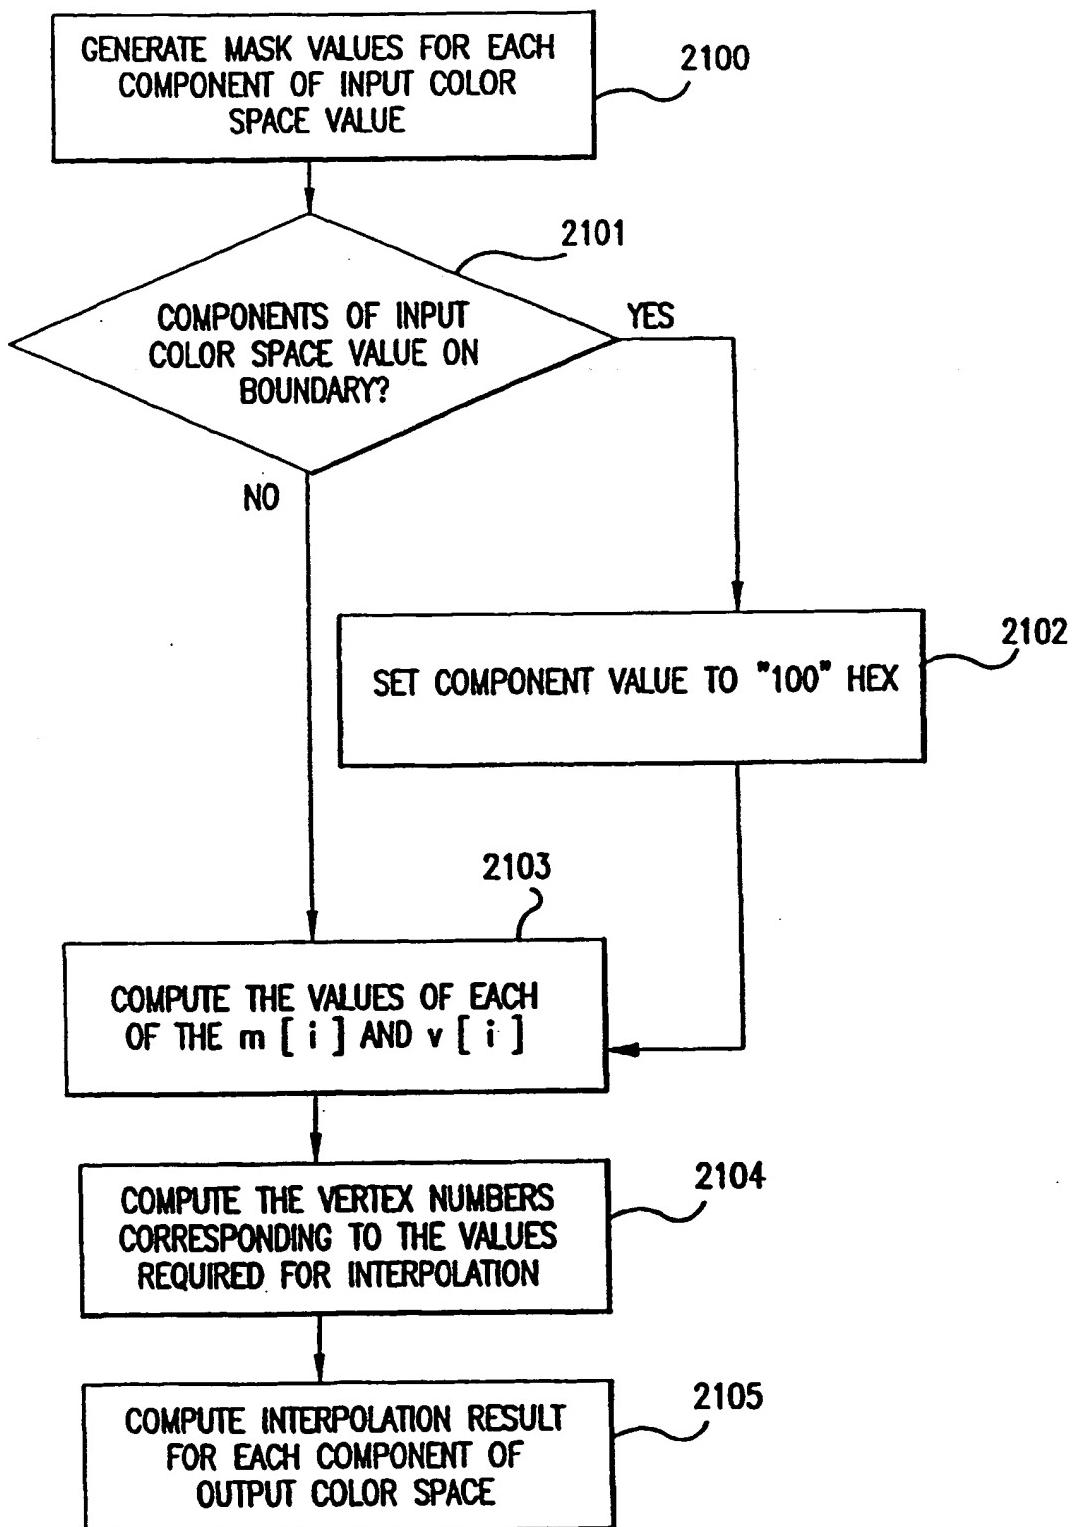


FIG.33

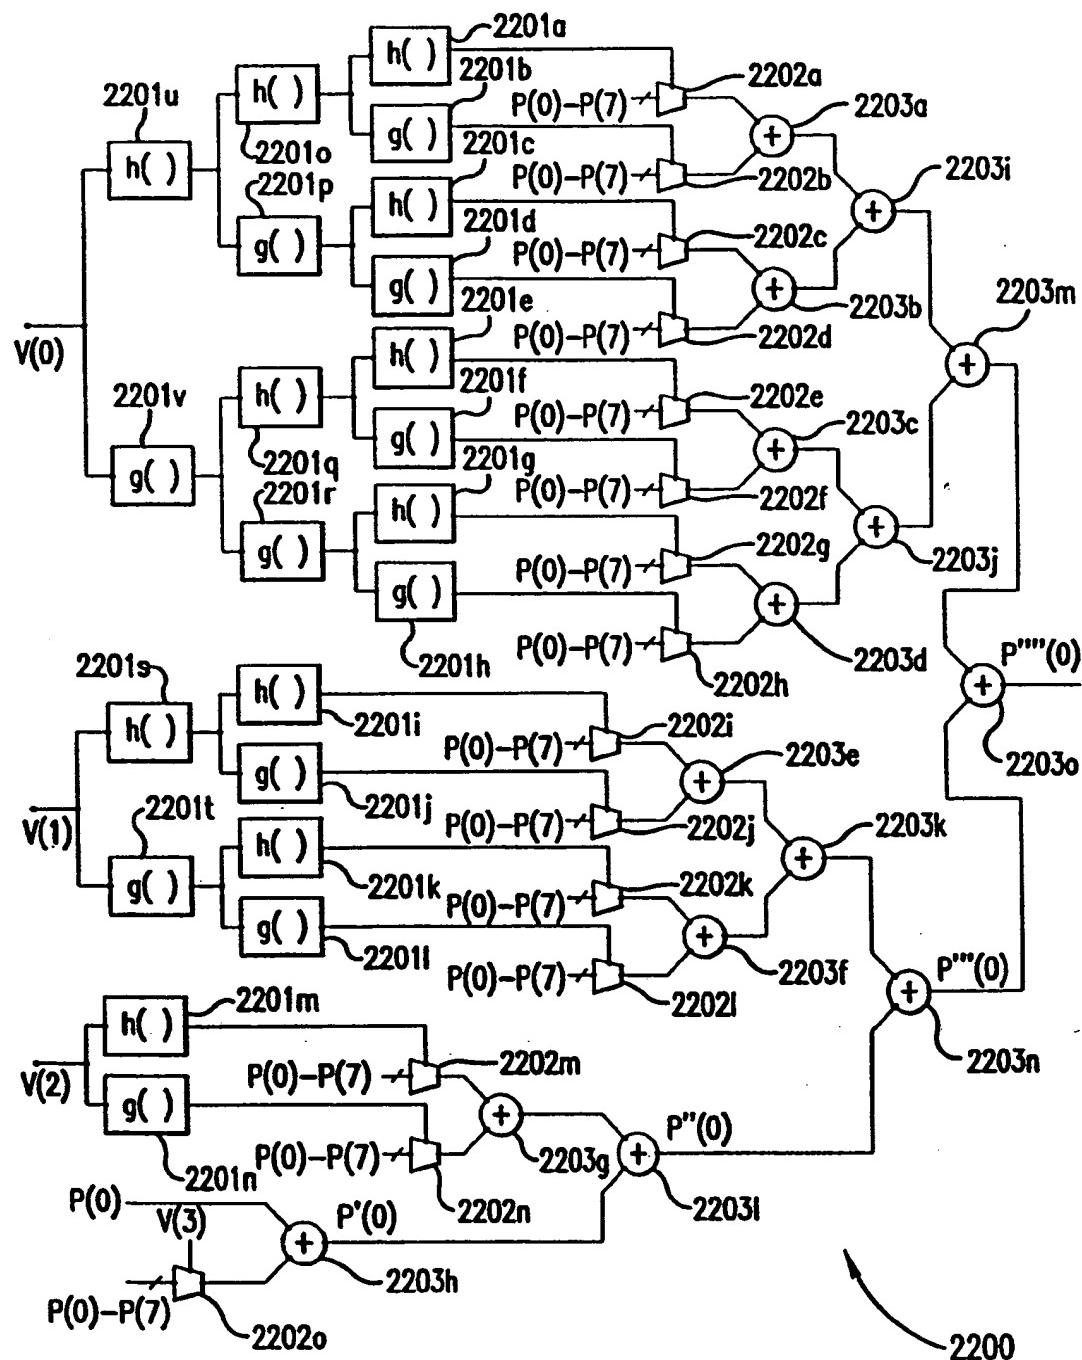


FIG. 34

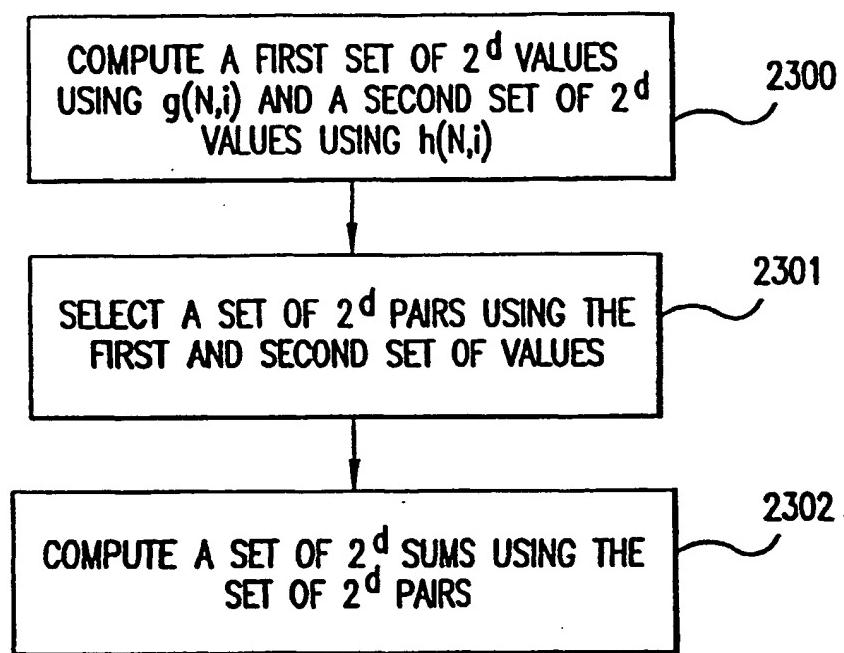
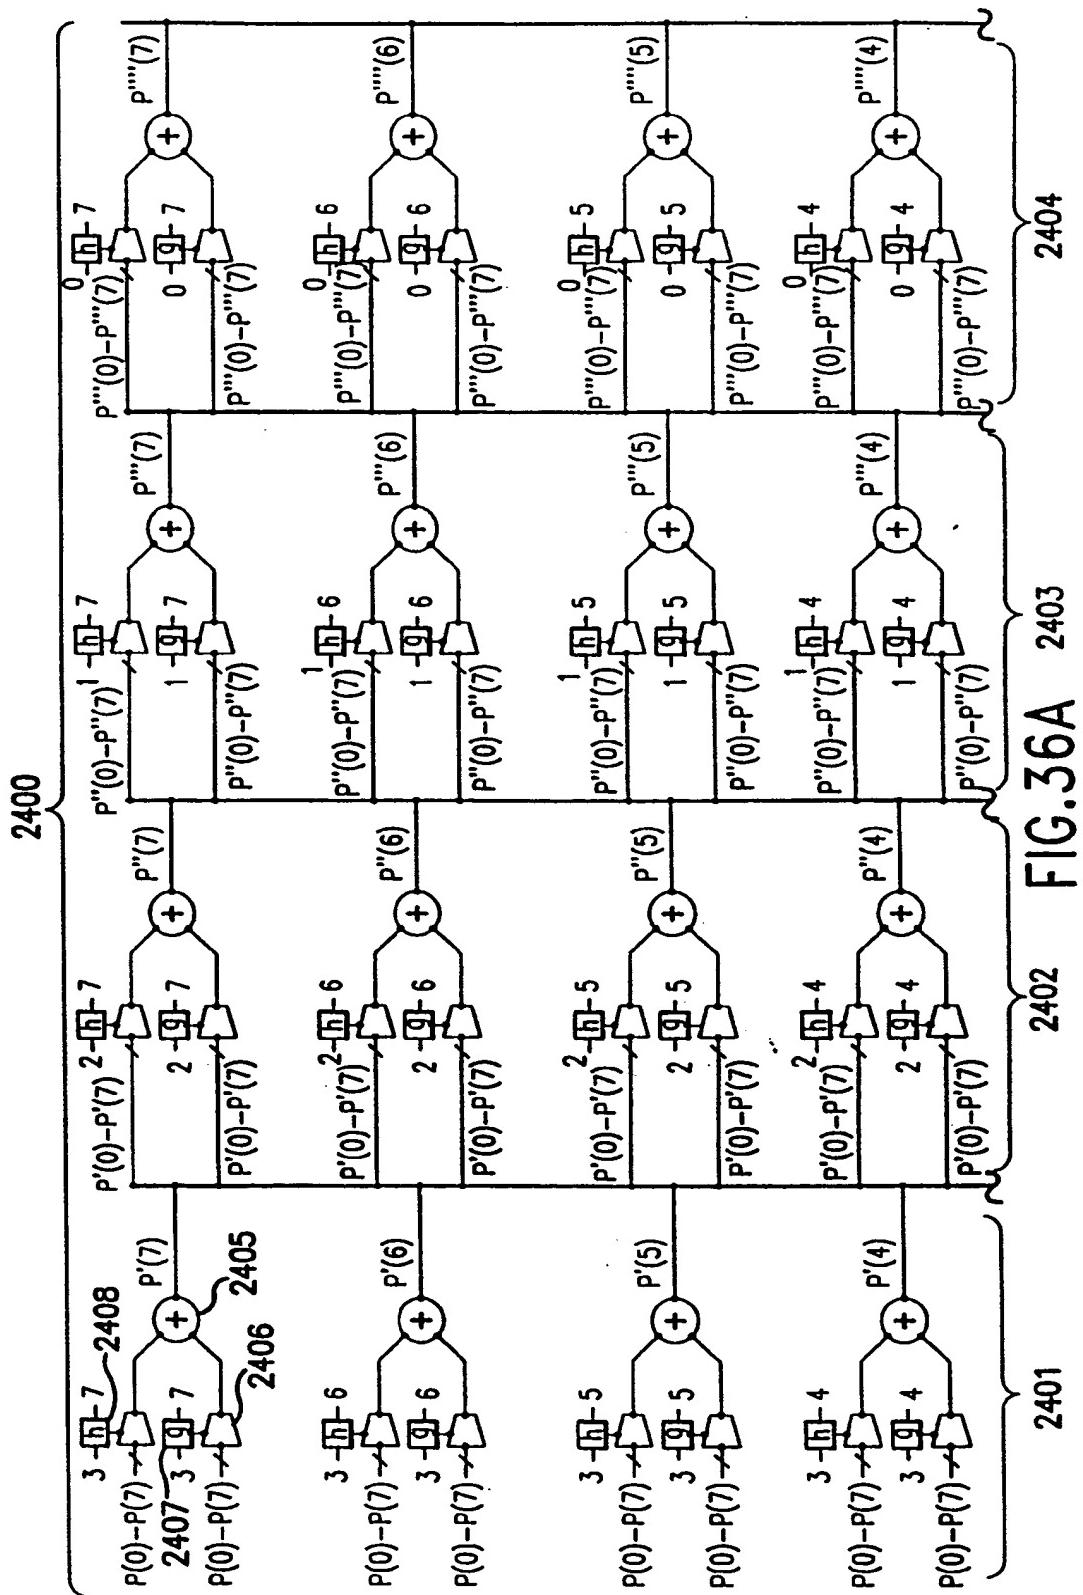
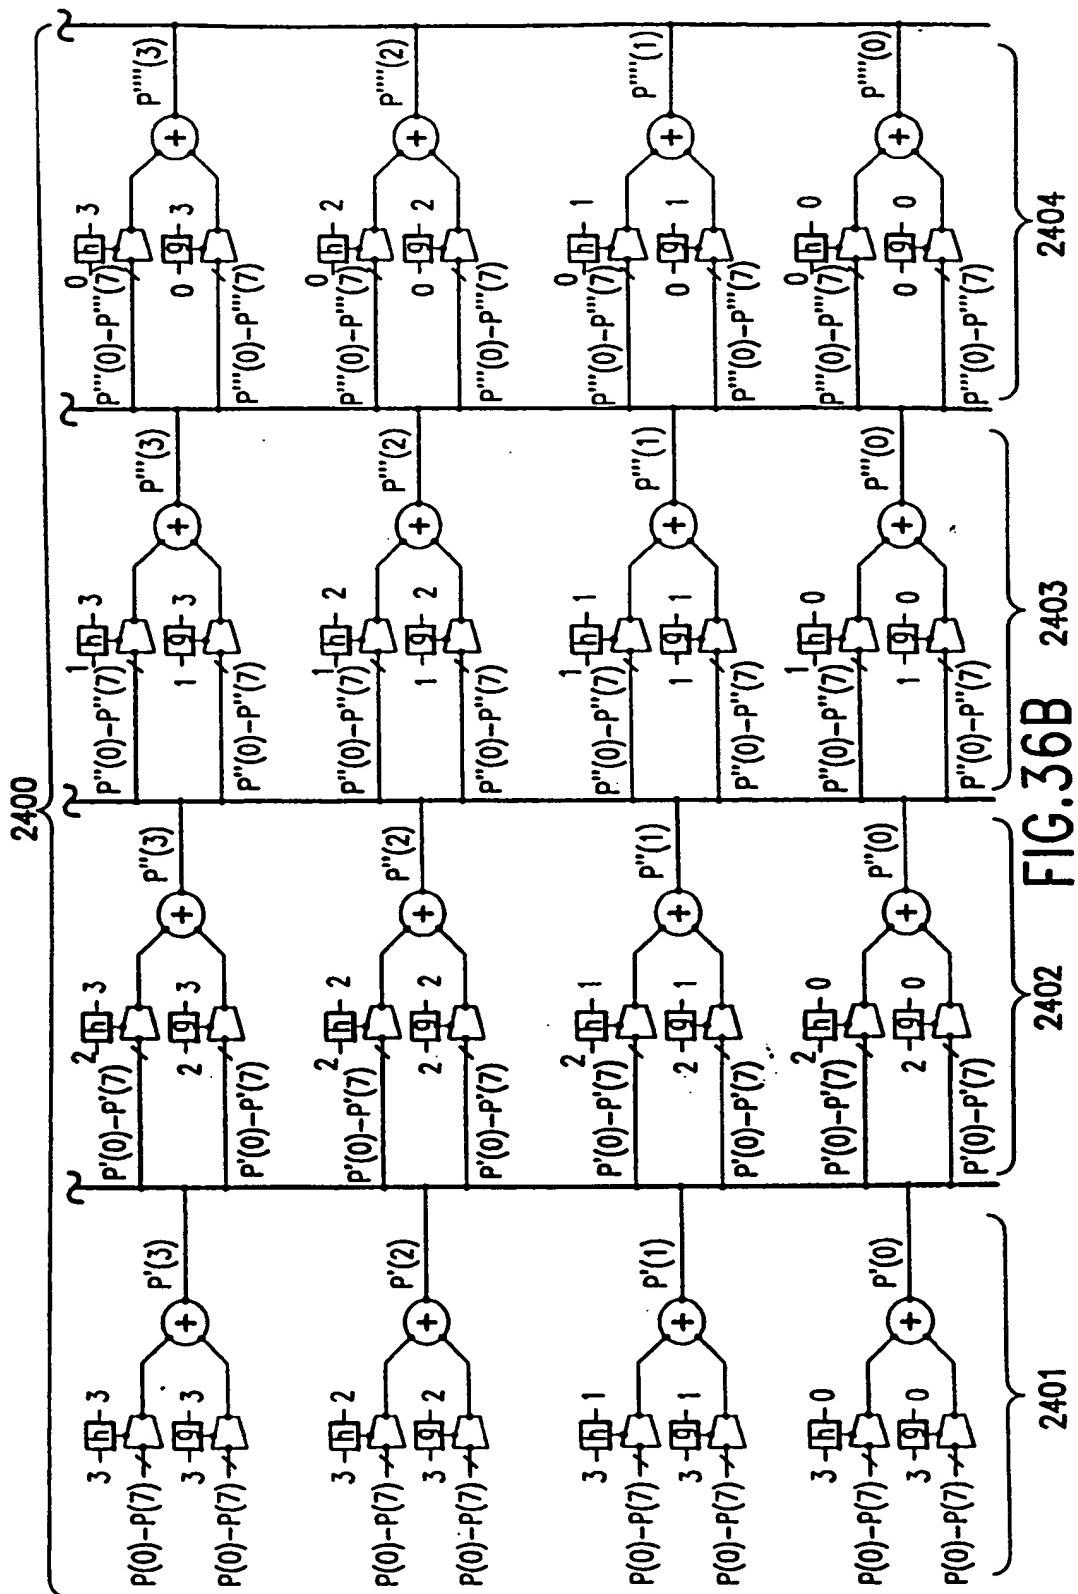


FIG.35





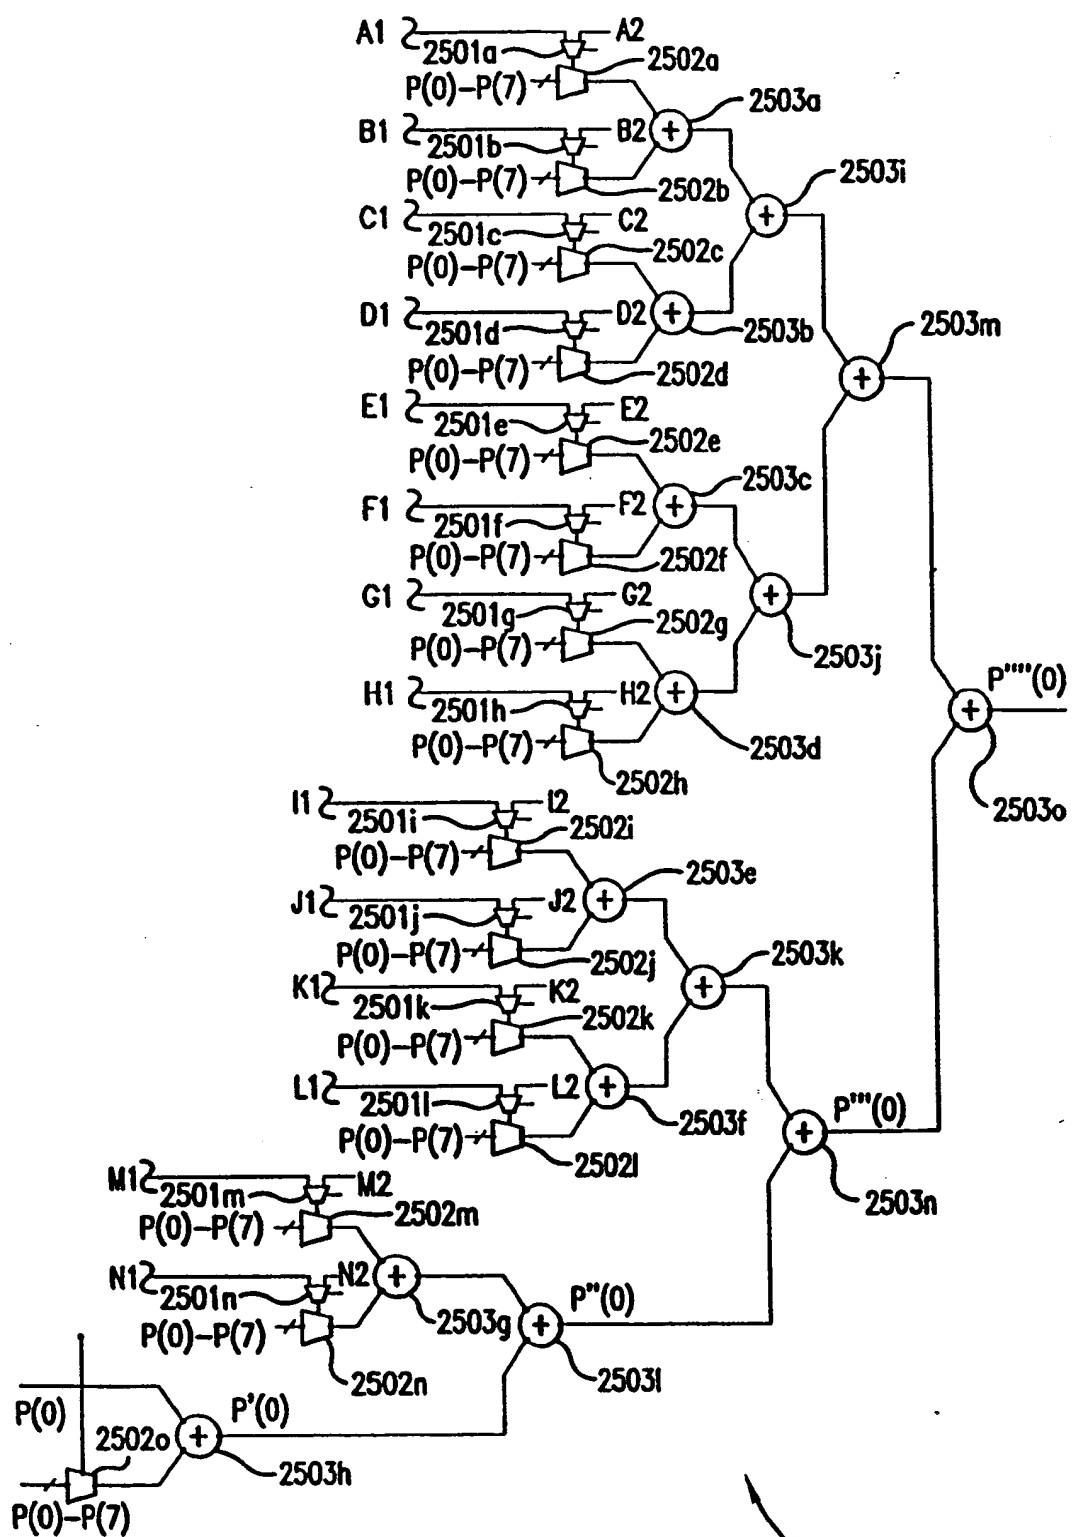
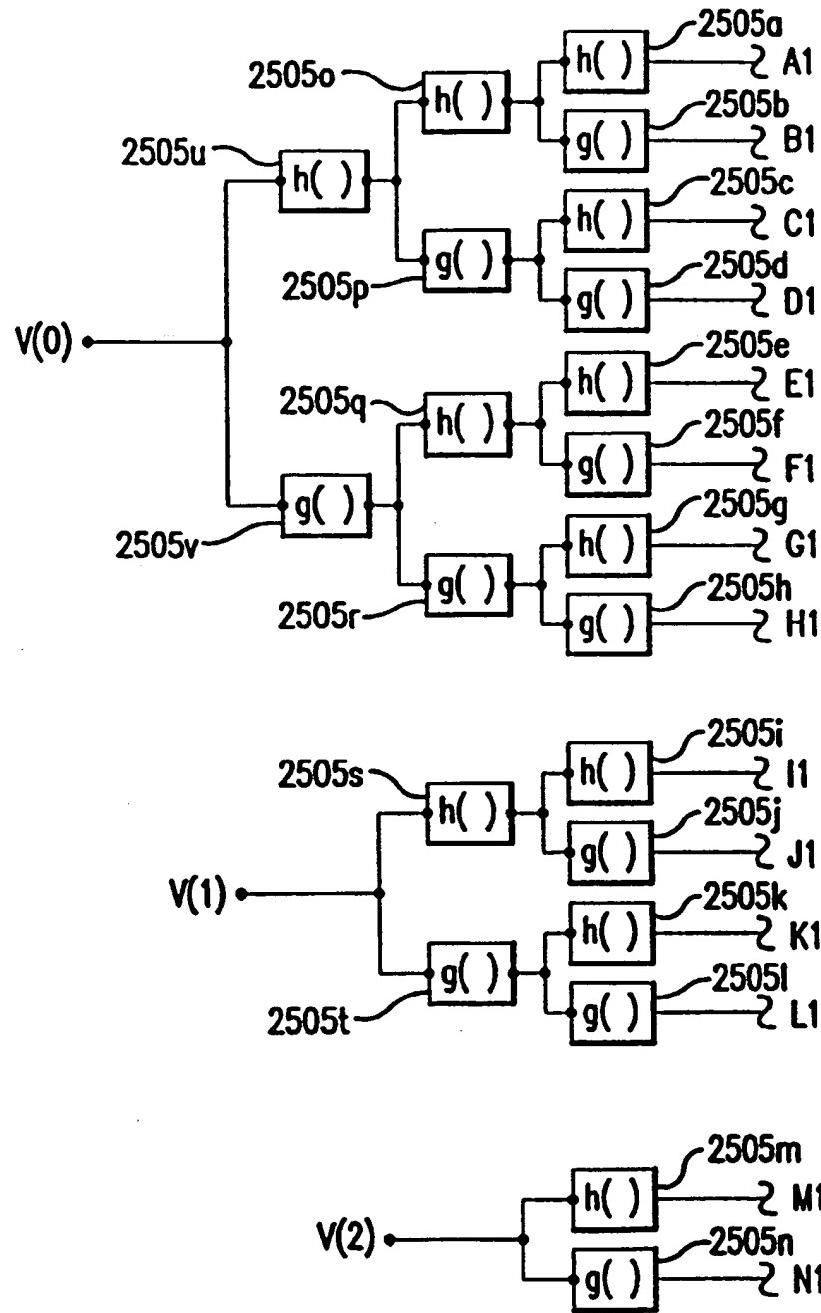


FIG. 37A



2500

FIG. 37B

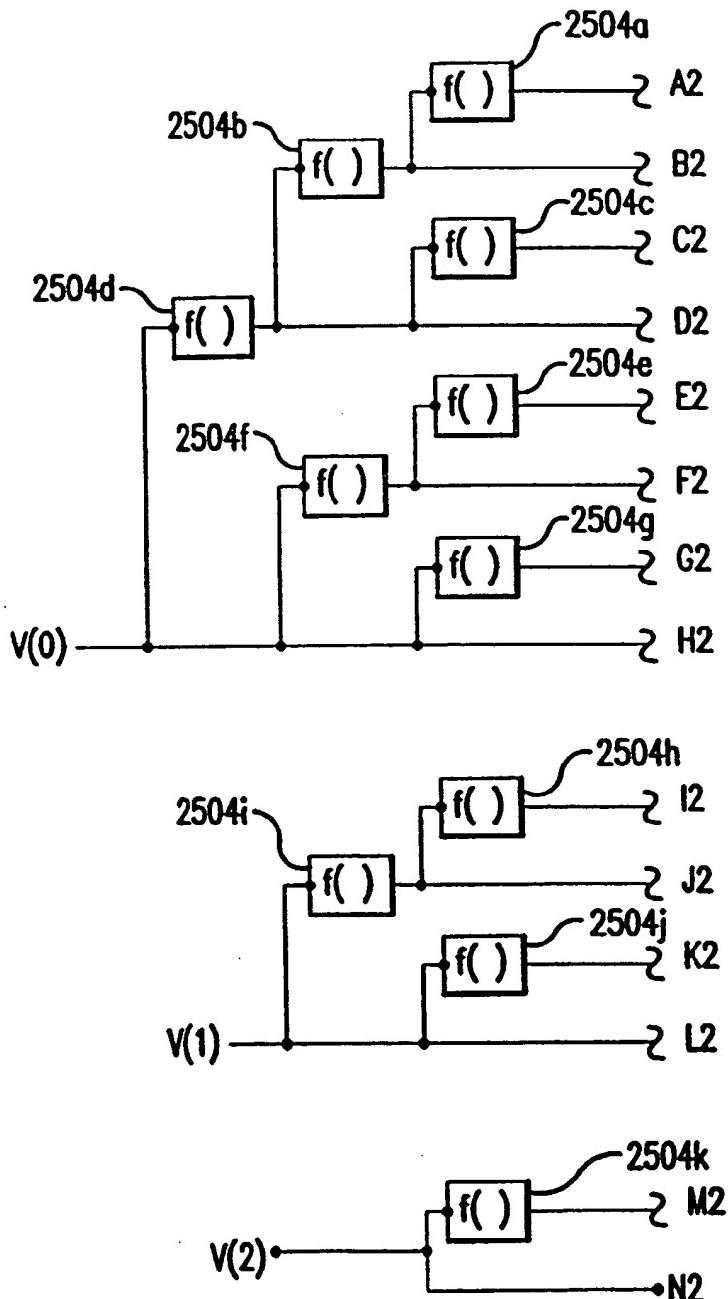


FIG. 37C

```

#define C_INC    1
#define B_INC   C_INC*17
#define A_INC   B_INC*17
#define TABLE_INC A_INC*17
/*
 * radial interpolation()
 */
void radial_interpolation(unsigned char input[], unsigned char output[])
{
    register int      origin, v0, v1, v2, v3;
    register int      a, b, c;
    register int      x, y, z;
    /*
     * Snap max value to last table value
     */
    a = (input[0] == 0xff) ? 0x100 : (int)input[0];
    b = (input[1] == 0xff) ? 0x100 : (int)input[1];
    c = (input[2] == 0xff) ? 0x100 : (int)input[2];
    /*
     * Compute Indexes into table
     */
    origin = A_INC*(a >> 4)
    + B_INC*(b >> 4)
    + (c >> 4);
    v0 = origin + A_INC*( (a & 0x01) )
    + B_INC*( (b & 0x01) )
    + (c & 0x01);
    v1 = origin + A_INC*( (a & 0x02) >> 1)
    + B_INC*( (b & 0x02) >> 1)
    + (c & 0x02) >> 1;
    v2 = origin + A_INC*( (a & 0x04) >> 2)
    + B_INC*( (b & 0x04) >> 2)
    + (c & 0x04) >> 2;
    v3 = origin + A_INC*( (a & 0x08) >> 3)
    + B_INC*( (b & 0x08) >> 3)
    + (c & 0x08) >> 3;
    /*
     * Compute x
     */
    x = ( ( (int) table[0][origin]
    + (int) table[0][v0]
    + (int) table[0][v1] << 1 )
    + ( (int) table[0][v2] << 2 )
    + ( (int) table[0][v3] << 3 )
    + 0x08 ) >> 4 ); /* Round */
    /*
     * Compute y
     */
    y = ( ( (int) table[1][origin]
    + (int) table[1][v0]
    + (int) table[1][v1] << 1 )
    + ( (int) table[1][v2] << 2 )
    + ( (int) table[1][v3] << 3 )
    + 0x08 ) >> 4 ); /* Round */
    /*
     * Compute z
     */
    z = ( ( (int) table[2][origin]
    + (int) table[2][v0]
    + (int) table[2][v1] << 1 )
    + ( (int) table[2][v2] << 2 )
    + ( (int) table[2][v3] << 3 )
    + 0x08 ) >> 4 ); /* Round */
    output[0] = (char)x;
    output[1] = (char)y;
    output[2] = (char)z;
}

```

FIG.38

```

-- Radial Interpolation.
entity radial is
port (
    -- Input Pixel (a,b,c) --> INPUT_ARRAY;
    -- Input: Vertex Points of Cube from Color Table --
    -- Point: POINT_ARRAY;
    -- Output Pixel (x,y,z) --> OUTPUT_ARRAY
);
end radial;

architecture radial1t of radial is
signal vertex: VERTEX_ARRAY;
signal point_value: POINT_VALUE_ARRAY;
begin
    VERTICES: process ( input )
begin
    -- Compute selected vertices from slices across input.
    -- For each slice across the input lower bits --
    -- Select the values of vertices to be used by --
    -- interpolation. The last two vertices are --
    -- always point(0) and the rounding constant. --
    for i in NUM_INTERP_BITS-1 downto 0 loop
        vertex(i) <= input(0)(i) & input(1)(i)
            & input(2)(i);
    end loop;
    end process;
    MUX: process ( vertex, point )
begin
    -- Initialize --
    point_value(d) <= (others => '0');
    -- point_value( 0 : NUM_INTERP_BITS-1 ) --
    for i in 0 to NUM_INTERP_BITS-1 loop
        case vertex(i) is
            when "000" => point_value(d)(i)(7+i downto i)
                when "001" => point_value(d)(i)(7+i downto i)
                when "010" => point_value(d)(i)(7+i downto i)
                when "011" => point_value(d)(i)(7+i downto i)
                when "100" => point_value(d)(i)(7+i downto i)
                when "101" => point_value(d)(i)(7+i downto i)
                when "110" => point_value(d)(i)(7+i downto i)
                when others => point_value(d)(i)(7+i downto i);
        end case;
    end loop;
    -- Always point(0) and the rounding constant. --
    -- For each slice across the input lower bits --
    -- Select the values of vertices to be used by --
    -- interpolation. The last two vertices are --
    -- always point(0) and the rounding constant. --
    for i in NUM_INTERP_BITS-1 downto 0 loop
        point_value(d)(i)(7+i downto i)
            <= point(d)(i);
    end loop;
end;

```

FIG.39A

```

-- Step 1: Move temp(even's) --
if ( 2*j < bound ) then
  temp(j) := temp(2*j) ;
else temp(j) := (others => '0') ;
end if ;
-- Rounding constant --
point_value(d)(NUM_INTERP_BITS)(7 downto 0) <= point(d)(0) ;
-- point_value(d)(NUM_INTERP_BITS+1)
-- (NUM_INTERP_BITS-1) <= '1' ;
end loop ;
end process ;
-- SUM adds the selected vertices values and truncates to eight bits. The process is implemented to minimize time by successive pairing and adding.
SUM: process ( point_value ) POINT_VALUES :
variable temp: INTEGER ;
variable bound;
begin
  -- For each output dimension --
  for d in 0 to NUM_OUTPUT_DIMEN-1 loop
    -- Initialize temp and bound --
    temp := point_value(d);
    bound := NUM_INTERP_BITS + 2 ;
    -- Loop for cell(log2(NUM_INTERP_BITS-1)) --
    for i in 0 to NUM_INTERP_BITS-1 loop
      -- Pair temp array elements and add --
      for j in 0 to (NUM_INTERP_BITS/2) loop
        -- Step 1: Move temp(even's) --
        if ( 2*j < bound ) then
          temp(j) := temp(2*j) ;
        else temp(j) := (others => '0') ;
        end if ;
        -- Step 2: Add temp(odd's) --
        if ( (2*j)+1 < bound ) then
          temp(j) := temp(j) + temp((2*j)+1) ;
        else temp(j) := temp(j) ;
        end if ;
        end loop ;
        -- Set bound for next iteration --
        bound := ( bound / 2 ) + ( bound mod 2 ) ;
      end loop ;
      -- Truncate output. --
      output(d) <= temp(0)(7+NUM_INTERP_BITS
                           downto NUM_INTERP_BITS) ;
    end loop ;
    end process ;
  end radialrt1;
end;

```

FIG.39B

FIG. 4.0A

EP 0 923 049 A2

```

+ index3 = index0 & 0x0001) :
index3 = index0
+ A_INC*{v3 and v2 and v1 & 0x0004} >> 2}
+ B_INC*{v3 and v2 and v1 & 0x0002} >> 1}
+ index4 = index0
+ A_INC*{v3 and v2 and v1 and v0 & 0x0004} >> 2}
+ B_INC*{v3 and v2 and v1 and v0 & 0x0002} >> 1}
+ index5 = index0
+ A_INC*{v3 and v2 and v1 or v0 & 0x0004} >> 2}
+ B_INC*{v3 and v2 and v1 or v0 & 0x0002} >> 1}
+ index6 = index0
+ A_INC*{v3 and v2 or v1 & 0x0004} >> 2}
+ B_INC*{v3 and v2 or v1 & 0x0002} >> 1}
+ index7 = index0
+ A_INC*{v3 and v2 or v1 and v0 & 0x0004} >> 2}
+ B_INC*{v3 and v2 or v1 and v0 & 0x0002} >> 1}
+ index8 = index0
+ A_INC*{v3 and v2 or v1 or v0 & 0x0004} >> 2}
+ B_INC*{v3 and v2 or v1 or v0 & 0x0002} >> 1}
+ index9 = index0
+ A_INC*{v3 or v2 & 0x0004} >> 2}
+ B_INC*{v3 or v2 & 0x0002} >> 1}
+ index10 = index0
+ A_INC*{v3 or v2 and v1 & 0x0004} >> 2}
+ B_INC*{v3 or v2 and v1 & 0x0002} >> 1}
+ index11 = index0
+ A_INC*{v3 or v2 and v1 and v0 & 0x0004} >> 2}
+ B_INC*{v3 or v2 and v1 and v0 & 0x0002} >> 1}
+ index12 = index0
+ A_INC*{v3 or v2 or v1 & 0x0004} >> 2}
+ B_INC*{v3 or v2 or v1 & 0x0002} >> 1}
+ index13 = index0
+ A_INC*{v3 or v2 or v1 and v0 & 0x0004} >> 2}
+ B_INC*{v3 or v2 or v1 and v0 & 0x0002} >> 1}
+ index14 = index0
+ A_INC*{v3 or v2 or v1 or v0 & 0x0004} >> 2}
+ B_INC*{v3 or v2 or v1 or v0 & 0x0002} >> 1}

/* Compute X */
x = ( (table[0][index0] + table[0][index1] + table[0][index2] + table[0][index3] + table[0][index4] + table[0][index5] + table[0][index6] + table[0][index7] + table[0][index8] + table[0][index9] + table[0][index10] + table[0][index11] + table[0][index12] + table[0][index13] + table[0][index14] + table[0][index15]) * Round*) / 0x08;
/* Compute Y */
y = ( (table[1][index0] + table[1][index1] + table[1][index2] + table[1][index3] + table[1][index4] + table[1][index5] + table[1][index6] + table[1][index7] + table[1][index8] + table[1][index9] + table[1][index10] + table[1][index11] + table[1][index12] + table[1][index13] + table[1][index14] + table[1][index15]) * Round*) / 0x08;
/* Compute Z */
z = ( (table[2][index0] + table[2][index1] + table[2][index2] + table[2][index3]) * Round*) / 0x08;

```

FIG.40B

```
+ table[2][index4] + table[2][index5]
+ table[2][index6] + table[2][index7]
+ table[2][index8] + table[2][index9]
+ table[2][index10] + table[2][index11]
+ table[2][index12] + table[2][index13]
+ table[2][index14] + table[2][index15]
+ 0x08 ) >> 4 ) ;
/*Round*/
*****
* Write back results */
*****
output[0] = (char) x ;
output[1] = (char) y ;
output[2] = (char) z ;
}
```

FIG.40C

```

#define f(v, m, N) (v | (N & -m))

/*
 * nonsym radial interpolation()
 */
void nonsym radial interpolation( unsigned char input[],
unsigned char point[], unsigned char resolution[],
unsigned char output[] )
{
    register unsigned char Amask, Bmask, Cmask ;
    register int a, b, c, x, y, z ;
    register unsigned char m0, m1, m2, m3 ;
    register unsigned char v0, v1, v2, v3 ;
    register unsigned char f_v1, v0_f, v2_f, v1_f ;
    register unsigned char f_v2_f, v1_v0_f, v2_v0_f ;
    register unsigned char f_v3_f, v2_v1_f, v3_v0_f ;
    register unsigned char f_v3_f, v2_f, v1_f, v0_f ;
    register unsigned char f_v3_f, v2_v0_f, v3_v1_f ;
    register unsigned char f_v3_f, v2_v1_f, v0_f, v1_f ;
    /* Generate Masks */
    Amask = ( 1 << resolution[0] ) - 1 ;
    Bmask = ( 1 << resolution[1] ) - 1 ;
    Cmask = ( 1 << resolution[2] ) - 1 ;
    /* Snap last value to last table value */
    a = ( input[0] == 0xff ) ? 0x100 : (int) input[0] ;
    b = ( input[1] == 0xff ) ? 0x100 : (int) input[1] ;
    c = ( input[2] == 0xff ) ? 0x100 : (int) input[2] ;
    /* Compute slices across masks and input components */
    m0 = ( Amask & 0x01 ) << 2 ) + ( Cmask & 0x01 );
    m1 = ( Amask & 0x02 ) << 1 ) + ( Cmask & 0x02 ) >> 1 );
    m2 = ( Amask & 0x04 ) + ( Bmask & 0x04 ) >> 1 ) + ( Cmask & 0x04 ) >> 2 );
    m3 = ( Amask & 0x08 ) >> 1 ) + ( Bmask & 0x08 ) >> 2 ) + ( Cmask & 0x08 ) >> 3 );
    /* Compute offset from origin. */
    /* Compute X */
    x = ( point[0][f_v1, v0_f, v2_f, v1_v0_f] +
    point[0][f_v2_f, v1_v0_f, v3_f, v2_v0_f] +
    point[0][f_v3_f, v2_v0_f, v3_f, v2_v1_f] +
    point[0][f_v3_f, v2_v1_f, v3_f, v2_v1_f] +
    point[0][f_v3_f, v2_v1_f, v0_f, v1_f] +
    point[0][f_v3_f, v2_v1_f, v0_f, v1_f, v0_f] ;
    /* Compute Y */
    y = ( point[0][f_v1, v0_f, v2_f, v1_v0_f] +
    point[0][f_v2_f, v1_v0_f, v3_f, v2_v0_f] +
    point[0][f_v3_f, v2_v0_f, v3_f, v2_v1_f] +
    point[0][f_v3_f, v2_v1_f, v3_f, v2_v1_f] +
    point[0][f_v3_f, v2_v1_f, v0_f, v1_f] +
    point[0][f_v3_f, v2_v1_f, v0_f, v1_f, v0_f] ;
    /* Compute Z */
    z = ( point[0][f_v1, v0_f, v2_f, v1_v0_f] +
    point[0][f_v2_f, v1_v0_f, v3_f, v2_v0_f] +
    point[0][f_v3_f, v2_v0_f, v3_f, v2_v1_f] +
    point[0][f_v3_f, v2_v1_f, v3_f, v2_v1_f] +
    point[0][f_v3_f, v2_v1_f, v0_f, v1_f] +
    point[0][f_v3_f, v2_v1_f, v0_f, v1_f, v0_f] ;
    /* Compute output */
    output[0] = ( m0 & ( (a & 0x01) << 2 ) +
    m1 & ( (b & 0x01) << 1 ) + ( Cmask & 0x01 ) );
    output[1] = ( m2 & ( (a & 0x02) << 1 ) + ( Bmask & 0x04 ) >> 1 ) +
    m3 & ( (a & 0x08) >> 2 ) + ( Cmask & 0x08 ) >> 3 );
}

```

FIG. 41A

FIG. 41 B

```

-- Pruned Tetrahedral Interpolation
entity pruned_tetrahedral is
port (
  -- Interpolation Select
  tetra_not_radial: in INPUT_ARRAY;
  -- Input Pixel ...
  input:          in INPUT_ARRAY;
  -- Vertex Points of Cube from Color Table ...
  point:         in POINT_ARRAY;
  -- Output Pixel ...
  output:        out OUTPUT_ARRAY
);
end pruned_tetrahedral;

architecture pruned_tetrahedral of pruned_tetrahedral is
begin
  constant VERTEX_COUNT: INTEGER := 2**NUM_INTERP_BITS;
  signal vertex:  VERTEX_ARRAY;
  signal point_value: POINT_VALUE_ARRAY;
  begin
    -- Compute vertices.
    VERTICES: process ( tetra_not_radial, input )
      variable vertex: STD_LOGIC_VECTOR(2 downto 0) ;
      variable temp1:  VERTEX_ARRAY;
      variable temp2:  VERTEX_ARRAY;
      variable limit, prev_limit: INTEGER;
    begin
      -- Initialize Array ...
      limit := 0;
      temp2 := ( others => "000" );
      -- For each slice across the interpolate bits ...
      for i in 0 to NUM_INTERP_BITS-1 loop
        -- Load slice into limit location ...
        vertex := input(0)(i) & input(1)(i)
        & input(2)(i);
        -- Select Radial or Tetrahedral --
        if ( tetra_not_radial = '0' ) then
          -- Radial ...
          -- limit is 1, 2, 4, 8 ...
          prev_limit := limit;
          limit := 2**i;
          -- Assign multiplex selects --
          for j in (VERTEX_COUNT/2)-1 downto 0
            loop
              if ( j < limit ) then
                temp2(prev_limit + j) := vertex;
              else
                temp2 := temp2;
              end if;
            end loop;
          else
            -- Tetrahedral ...
            -- limit is 0, 2, 6, 14, 30, ...
            limit := ( 2**i+1 ) - 2;
            -- Load slice into limit location ...
          end if;
        end if;
      end loop;
    end process;
  end;
end;

```

FIG.42A

```

temp2(1limit) := vertex ;
temp1 := temp2 ;
-- for all values in temp1 --
for j in (VERTEX_COUNT/2)-1 downto 0
loop if ( j < limit/2 ) then
      temp2(2*j) := temp1(1limit)
      and temp1(j) := temp2(j)
      temp2((2*j)+1) := temp1(j)
      or temp1(j) ;
else temp2 := temp2 ;
      end if ;
      end loop ;
end loop ;
-- Move into index array --
vertex <- temp2 ;
end process ;
Select the values of vertices to be used by the
-- interpolation. The last two vertices are always
-- point(0) and the rounding constant.
MUX: process ( vertex, point )
begin for d in 0 to NUM_OUTPUT_DIMEN-1 loop
point_value(d) <- others => '0') ;
for i in 0 to VERTEX_COUNT-2 loop
case vertex(i) is
when "000" => point_value(d)(1)(7 downto 0) :=
point_value(d)(1)(7 downto 0) ;
when "010" => point_value(d)(1)(7 downto 0) :=
point_value(d)(1)(7 downto 0) ;
when "011" => point_value(d)(1)(7 downto 0) :=
point_value(d)(1)(7 downto 0) ;
when "100" => point_value(d)(1)(7 downto 0) :=
point_value(d)(1)(7 downto 0) ;
when "101" => point_value(d)(1)(7 downto 0) :=
point(d)(5) ;
when "110" => point_value(d)(1)(7 downto 0) :=
point(d)(6) ;
when others => point_value(d)(1)(7 downto 0) :=
point(d)(7) ;
end case ;
end loop ;
point_value(d)(VERTEX_COUNT-1) :=
point(d)(0) ;
point_value(d)(VERTEX_COUNT) :=
point(d)(0) <- point(d)(1) ;
end loop ;
end process ;
-- SUM adds the selected vertices values and
-- truncates to eight bits. The process is
-- implemented to minimize prop delay by
-- successive pairing and adding.
SUM: process ( point_value )
variable temp: POINT_VALUES ;
variable bound: INTEGER ;
begin for d in 0 to NUM_OUTPUT_DIMEN-1 loop
temp := point_value(d) ;
point_value(d) := point_value(d) +1 ;
for j in 0 to NUM_INTERP_BITS loop
for i in 0 to VERTEX_COUNT loop
if ( 2**j < bound ) then
      temp(j) := temp(2**j) ;
else temp(j) := (others => '0') ;
end if ;
temp(j) := (others => '0') ;
if ( (2**j)+1 < bound ) then
      temp(j) := temp(2**j) ;
end if ;
end loop ;
end loop ;
end process ;

```

FIG. 42B

```
temp(j) := temp(j) + temp((2*j)+1) ;
else
    temp(j) := temp(j) ;
end if ;
end loop ;
bound := ( bound / 2 ) + ( bound mod 2 ) ;
end loop ;
output(d) <- temp(0)(7+NUM_INTERP_BITS
downto NUM_INTERP_BITS) ;
end loop ;
end process ;
-----;
end pruned_tetrtl ;
-----;
```

FIG.42C

```

-- Pruned Tetrahedral Interpolation
entity pruned_tetrahedral is
port (
    -- Input Pixel --
    input:     in INPUT_ARRAY;
    -- Vertex Points of Cube from Color Table --
    point:    in POINT_ARRAY;
    -- Output Pixel --
    output:   out OUTPUT_ARRAY
);
end pruned_tetrahedral;

architecture pruned_tetrahedral of pruned_tetrahedral is
constant VERTEX_COUNT: INTEGER := 2**NUM_INTERP_BITS;
signal vertex:  VERTEX_ARRAY;
signal point_value: POINT_VALUE_ARRAY;
begin
    -- Compute vertices --
    VERTICES: process ( input )
        variable temp1:  VERTEX_ARRAY : INTEGER;
        variable limit: INTEGER;
    begin
        -- Initialize Array --
        temp2 := ( others => "000" );
        -- For each slice across the interpolate bits --
        for i in 0 to NUM_INTERP_BITS-1 loop
            -- limit is 0, 2, 6, 14, 30, ...
            -- limit := ( 2*(i+1) ) - 2 ;
            -- Load slice into limit location --
            temp2(1limit) := input(0)(i) & input(1)(i) & input(2)(i);
            temp1 := temp2 ;
            -- for all values in temp1 --
            for j in (VERTEX_COUNT/2)-1 downto 0 loop
                if ( j < 1limit/2 ) then
                    temp2(2*j) := temp1(j);
                    temp2((2*j)+1) := temp1(j);
                    or temp1(j);
                else
                    temp2 := temp2 ;
                end if ;
            end loop ;
            -- Move into index array --
            vertex <- temp2 ;
        end process ;
    end;
    -- Select the values of vertices to be used by the --
    -- interpolation. The last two vertices are always --
    -- point(0) and the rounding constant.
    NUX: process ( vertex, point )
begin
    -- For each output dimension --
    for d in 0 to NUM_OUTPUT_DIMEN-1 loop
        -- Initialize --
        point_value(d) <- (others => '0' );
    end loop;
end;

```

FIG. 43A

```

-- For each output dimension --
for d in 0 to NUM_OUTPUT_DIMEN-1 loop
  -- Initialize temp and bound. --
  temp := point_value(d);
  bound := VERTEX_COUNT + 1;
  -- Loop for ceil(log2(VERTEX_COUNT)) --
  for i in 0 to NUM_INTERP_BITS loop
    -- Pair temp array elements and add. --
    for j in 0 to VERTEX_COUNT/2 loop
      -- Step 1: Move temp(even's) --
      if ( 2*j < bound ) then
        temp(j) := temp(2*j);
      else
        temp(j) := (others => '0');
      end if;
    -- Step 2: Add temp(odd's) --
    if ( (2*j)+1 < bound ) then
      temp(j) := temp(j) + temp((2*j)+1);
    else
      temp(j) := temp(j);
    end if;
  end loop;
  -- point(d)(0) --
  point_value(d)(VERTEX_COUNT-1) <- point(d)(0);
  -- Rounding constant --
  point_value(d)(VERTEX_COUNT-(NUM_INTERP_BITS-1)) <- '1';
end loop;
end process;
-- SUM adds the selected vertices values and --
-- truncates to eight bits. The process is --
-- implemented to minimize prop delay by --
-- successive pairing and adding.
SUM: process ( point_value )
variable temp: POINT_VALUES;
variable bound: INTEGER;
begin

```

FIG.43B

```
end loop ;  
-- Set bound for next iteration --  
bound := ( bound / 2 ) + ( bound mod 2 ) ;  
end loop ;  
-- Truncate output. --  
output(d) <= temp(0)(7+NUM_INTERP_BITS  
downto NUM_INTERP_BITS) ;  
end loop ;  
end process ;  
end pruned_tetrt1 ;
```

FIG.43C

```

top := top ;
    end if ;
end loop ;
-- Generate Mask --
for j in NUM_INTERP_BITS-1 downto 0 loop
    if ( j > top ) then
        mask(j)(1) <= '1' ;
    else
        mask(j)(1) <= '0' ;
    end if ;
end loop ;
end loop ;
end process ;
-- Compute Vertices --
VERTICES: Process ( input, mask )
variable temp1, temp2: VERTEX_ARRAY ;
variable limit: INTEGER ;
begin
    -- Initialize Arrays --
    temp2 := ( others => "000" ) ;
    -- For each slice across the interpolate bits --
    for i in 0 to NUM_INTERP_BITS-1 loop
        -- limit is 0, 2, 6, 14, 30, ... --
        limit := ( 2**((i+1)) ) - 2 ;
        -- Load slice into limit location ...

```

-- Non-Symmetric Radial Interpolation --

entity nonsym_radial is

port (
 -- Input Pixel --
 input: INOUT Pixel ;
 -- Vertex Points of Cube from Color Table --
 point: INOUT POINT_ARRAY ;
 -- Lattice Resolution --
 resolution: INOUT RESOLUTION_ARRAY ;
 -- Output Pixel --
 output: OUTOUT_PIXEL ;
 -- Output Array --
 output: OUTOUT_OUTPUT_ARRAY
)

end nonsym_radial ;

architecture nonsym_radial of nonsym_radial is

constant VERTEX_COUNT: INTEGER := 2**NUM_INTERP_BITS ;
signal mask: MASK_ARRAY ;
signal vertex: VERTEX_ARRAY ;
signal point_value: POINT_VALUE_ARRAY ;
begin

-- Compute Lattice Resolution Masks --
MASKS: process (resolution) INTEGER : variable top: begin

-- For each Component --
for i in 0 to 2 loop
 top := 0 ;
 -- Convert BIT_VECTOR to INTEGER --
 for j in 2 downto 0 loop
 if (resolution(i)(j) = '1') then
 top := top + 2**j ;
 else

FIG.44A

```

temp2(limit) := mask(i) and ( input(0)(i) & input(2)(i) ) :           <= point(d)(0)
temp1 := temp2 :           <= point(d)(1)
for all values in templ --           <= point(d)(2)
for j in (VERTEX_COUNT/2)-1 downto 0 loop
  if ( j < 11mit/2 ) then
    P[2*j] = P[j] ...
    temp2(2*j) := templ(j) :
    P[2*j+1]=v[i](P[j] & ~m[i]) :
  end if :
  temp2((2*j)+1) := templ(limit)
  or ( templ(j) and not mask[i] ) :
else
  temp2 := temp2 :
end if :
end loop :
end loop :
vertex <= temp2 :
vertex ( vertex, point )
begin
  for d in 0 to NUM_OUTPUT_DIMEN-1 loop
    point_value(d) := (others => '0');
    for i in 0 to VERTEX_COUNT-2 loop
      case vertex(i) is
        when "000" => point_value(d)(1)(7 downto 0)
        when "001" => point_value(d)(1)(7 downto 0)
        when "010" => point_value(d)(1)(7 downto 0)
        when "011" => point_value(d)(1)(7 downto 0)
        when "100" => point_value(d)(1)(7 downto 0)
        when "101" => point_value(d)(1)(7 downto 0)
        when "110" => point_value(d)(1)(7 downto 0)
        when others => point_value(d)(1)(7 downto 0)
      end case :
    end loop :
    point_value(d)(VERTEX_COUNT-1) := point(d)(0)
    point_value(d)(VERTEX_COUNT) := point(d)(6)
    point_value(d)(VERTEX_COUNT-1) <= '1'
  end loop :
end process :
----- Select the values of vertices to be used by the
----- interpolation. The last two vertices are always
----- point(0) and the rounding constant.
SUM: process ( point_value ) variable temp: POINT_VALUES ;
begin
  for d in 0 to NUM_OUTPUT_DIMEN-1 loop
    temp := point_value(d);
    bound := VERTEX_COUNT + 1;
    for i in 0 to NUM_INTERP_BITS loop
      for j in 0 to VERTEX_COUNT/2 loop
        if ( 2*j < bound ) then
          MUX: process ( vertex, point )
            begin
              for d in 0 to NUM_OUTPUT_DIMEN-1 loop
                point_value(d) := (others => '0');
                for i in 0 to VERTEX_COUNT-2 loop
                  case vertex(i) is
                    when "000" => point_value(d)(1)(7 downto 0)

```

FIG.44B

```
temp(j) := temp(2*j) ;
else      temp(j) := (others => '0') ;
end if ;
if ( (2*j)+1 < bound ) then
  temp(j) := temp(j) + temp((2*j)+1) ;
else
  temp(j) := temp(j) ;
end if ;
end loop ;
bound := ( bound / 2 ) + ( bound mod 2 ) ;
end loop ;
output(d) <= temp(0)(7+NUM_INTERP_BITS
downto NUM_INTERP_BITS) ;
end loop ;
end process ;
-----  
end nonsym_radialrt1 ;  
-----
```

FIG.44C

FIG. 45A

FIG. 45B

```

-- Non-Symmetric Pruned Tetrahedral Interpolation --
entity nonsym_pruned_tetrahedral is
port (
  -- Input Pixel --
  input:    in  INPUT_ARRAY ;
  -- Vertex Points of Cube from Color Table --
  point:   in  POINT_ARRAY ;
  -- Lattice Resolution --
  resolution: in  RESOLUTION_ARRAY ;
  -- Output Pixel --
  output:  out  OUTPUT_ARRAY
);
end nonsym_pruned_tetrahedral;

architecture nonsym_tetrtl of nonsym_pruned_tetrahedral is
begin
  constant VERTICES_COUNT: INTEGER := 2+NUM_INTERP_BITS;
  signal mask: MASK_ARRAY;
  signal vertex: VERTEX_ARRAY;
  signal point_value: POINT_VALUE_ARRAY;
  begin
    -- Compute Lattice Resolution Masks --
    MASKS: process ( resolution )  INTEGER : top;
    begin
      -- For each Component --
      for i in 0 to 2 loop
        top := 0;
        -- Convert BIT_VECTOR to INTEGER --
        for j in 2 downto 0 loop
          if ( resolution(i)(j) = '1' ) then
            top := top + 2**j;
          else
            top := top - 2**j;
          end if;
        end loop;
      end process;
    end if;
    -- Generate Mask --
    for j in NUM_INTERP_BITS-1 downto 0 loop
      if ( j <= top ) then
        mask(j)(i) <= '1';
      else
        mask(j)(i) <= '0';
      end if;
    end loop;
  end if;
end loop;
end process;
end;
-- Compute Vertices --
VERTICES: process ( input, mask )
variable temp1: TEMP1; temp2: VERTEX_ARRAY;
variable limit: INTEGER;
begin
  -- Initialize Arrays --
  temp2 := ( others => "000" );
  -- For each slice across the interpolate bits --
  for i in 0 to NUM_INTERP_BITS-1 loop
    -- limit is 0, 2, 6, 14, 30, ...
    limit := ( 2**i+1 ) - 2;
    -- Compute Vertices --
    for j in 0 to 2 loop
      vertex(j) := temp1;
      temp1 := temp2;
    end loop;
  end loop;
end;

```

FIG. 46A

```

-- Load slice into limit location
temp2(limit) := mask(i) and {input(2)(i)}' ;
temp1 := temp2;
-- for all values in templ --
for j in (VERTEX_COUNT/2)-1 downto 0 loop
  if (j < limit/2) then
    P'[2*j] := (v[i] | ~m[i]) & P[j]--;
    temp2(2*j) := (temp1[limit]
      or not mask(i) and temp1(j)) :
    -- P'[2*j+1] = v[i] | P[j]--;
    temp2((2*j)+1) := temp1[limit]
      or temp1(j) :
  else
    temp2 := temp2 ;
  end if;
end loop;
-- Move into index array --
vertex <- temp2;
end process;
-- Select the values of vertices to be used by the --
-- interpolation. The last two vertices are always --
-- point(0) and the rounding constant.
-- MUX: process ( vertex, point )
begin
  for d in 0 to NUM_OUTPUT_DIMEN-1 loop
    point_value(d) <- (others => '0');
    for i in 0 to VERTEX_COUNT-2 loop
      case vertex(i) is
        when "000" => point_value(d)(i)(7 downto 0)
          <- point(d)(0);
        when "001" => point_value(d)(i)(7 downto 0)
          <- point(d)(1);
        when "010" => point_value(d)(i)(7 downto 0)
          <- point(d)(2);
        when "011" => point_value(d)(i)(7 downto 0)
          <- point(d)(3);
        when "100" => point_value(d)(i)(7 downto 0)
          <- point(d)(4);
        when "101" => point_value(d)(i)(7 downto 0)
          <- point(d)(5);
        when "110" => point_value(d)(i)(7 downto 0)
          <- point(d)(6);
        when others => point_value(d)(i)(7 downto 0)
          <- point(d)(7);
      end case;
    end loop;
    point_value(d)(VERTEX_COUNT-1)
      <- point(d)(0);
    point_value(d)(VERTEX_COUNT)
      <- point(d)(6);
    (NUM_INTERP_BITS-1) <- '1';
  end loop;
end process;
----- Select the values of vertices to be used by the --
----- interpolation. The last two vertices are always --
----- point(0) and the rounding constant.
----- SUM: process ( point_value )
variable temp;
variable bound;
begin
  for d in 0 to NUM_OUTPUT_DIMEN-1 loop
    temp := point_value(d);
    bound := VERTEX_COUNT + 1;
    for i in 0 to NUM_INTERP_BITS-1 loop
      for j in 0 to VERTEX_COUNT/2 loop

```

FIG. 46B

```

if ( 2*j < bound ) then
temp(j) := temp(2*j) ;
else
temp(j) := (others => '0') ;
end if ;
if ( (2*j)+1 < bound ) then
temp(j) := temp(j) + temp((2*j)+1) ;
else
temp(j) := temp(j) ;
end if ;
end loop ;
bound := ( bound / 2 ) + ( bound mod 2 ) ;
end loop ;
output(d) <- temp(0)(7+NUM_INTERP_BITS
downto NUM_INTERP_BITS) ;
end loop ;
end process ;

```

.....
.....

```

end nonsym_pruned_tetrtl ;

```

FIG.46C

```

if ( j < top ) then
  mask(j)(i) <= '1'
else
  mask(j)(i) <= '0'
end if ;
end loop ;
end loop ;
end process ;
-- Compute vertices
-- VERTICES: process ( tetra_not_radial, input, mask )
variable temp1: VERTEX_ARRAY;
variable temp2: INTEGER;
begin
temp2 := ( others => "000" );
for i in 0 to NUM_INTERP_BITS-1 loop
  limit := ( 2**i+1 ) - 2 ;
  temp2(i) := mask(i) and ( input(0)(1)
                            & input(1)(1) & input(2)(1) );
  temp1 := temp2;
  for j in (VERTEX_COUNT/2)-1 downto 0 loop
    if ( j < limit/2 ) then
      -- Select Radial or Tetrahedral1 ...
      if ( tetra_not_radial = '0' ) then
        ... Radial ...
      else
        ... P'[2*j] = P[j] ...
        ... temp2(2*j) := temp1(j) ...
        ... P'[2*j+1]=v[1](P[j]&-m[i])
    end if ;
  end loop ;
  for j in NUM_INTERP_BITS-1 downto 0 loop
    if ( resolution(i)(j) = '1' ) then
      top := top + 2*j;
    else
      top := top ;
    end if ;
  end loop ;
end process;
-- Compute Lattice Resolution Masks
MASKS: process ( resolution ) INTEGER :
begin
for i in 0 to 2 loop
  top := 0 ;
  for j in 2 downto 0 loop
    if ( resolution(i)(j) = '1' ) then
      top := top + 2*j;
    else
      top := top ;
    end if ;
  end loop ;
end process;

```

FIG. 47A

```

temp2((2*j)+1) := temp1(limt)
or (temp1(j) and not mask(i));
else
  -- NonSym Pruned Tetrahedra
  -- P'[2*j] = (v[i])&P[j] --
  temp2(2*j) := ( temp1(limt)
  or not mask(i) ) and temp1(j);
  -- P'[2*j+1] = v[1]|P[j] --
  temp2((2*j)+1) := temp1(limt)
  or temp1(j);
end if;
else
  temp2 := temp2;
end if;
end loop;
end loop;
vertex <= temp2;
end process;
-- Select the values of vertices to be used by the
-- interpolation. The last two vertices are always
-- point(0) and the rounding constant.
MUX: process ( vertex, point )
begin
  for d in 0 to NUM_OUTPUT_DIMEN-1 loop
    point_value(d) <= (others => '0');
    for i in 0 to VERTEX_COUNT-2 loop
      case vertex(i) is
        point_value(d)(i)(7 downto 0) <=
          point(d)(0);
        when "001" => point_value(d)(i)(7 downto 0) <=
          point(d)(1);
        when "010" => point_value(d)(i)(7 downto 0) <=
          point(d)(2);
        when "011" => point_value(d)(i)(7 downto 0) <=
          point(d)(3);
        when "100" => point_value(d)(i)(7 downto 0) <=
          point(d)(4);
        when "101" => point_value(d)(i)(7 downto 0) <=
          point(d)(5);
        when "110" => point_value(d)(i)(7 downto 0) <=
          point(d)(6);
        when others => point_value(d)(i)(7 downto 0) <=
          point(d)(7);
      end case;
    end loop;
    point_value(d)(VERTEX_COUNT-1) <=
      point(d)(0);
    point_value(d)(VERTEX_COUNT) <=
      point(d)(NUM_INTERP_BITS-1) <= '1';
  end loop;
end process;
-- Select the values of vertices to be used by the
-- interpolation. The last two vertices are always
-- point(0) and the rounding constant.
SUM: process ( point_value )
variable temp;
variable bound;
begin
  for d in 0 to NUM_OUTPUT_DIMEN-1 loop
    temp := point_value(d);
    bound := VERTEX_COUNT + 1;
    for i in 0 to NUM_INTERP_BITS loop
      for j in 0 to VERTEX_COUNT/2 loop
        if ( 2*j < bound ) then
          temp(j) := temp(2*j);
        else
          temp(j) := (others => '0');
        end if;
      end loop;
    end loop;
  end loop;
end process;

```

FIG. 47B

```
end if ;
if ( (2*j)+1 < bound ) then
  temp(j) := temp(j) + temp((2*j)+1)
else
  temp(j) := temp(j) ;
end if ;
end loop ;
bound := ( bound / 2 ) + ( bound mod 2 ) ;
end loop ;
output(d) <= temp(0)(7+NUM_INTERP_BITS
                      downto NUM_INTERP_BITS) ;
end loop ;
end process ;
end nonsym_radial_tetrl ;
```

FIG.47C

THIS PAGE BLANK (USPTO)